Provably Approximated Point Cloud Registration

Ibrahim Jubran[1], Alaa Maalouf[1], Ron Kimmel[2], Dan Feldman[1]
[1]University of Haifa, [2]Technion
{ibrahim.jub, alaamalouf12, dannyf.post}@gmail.com
{ron}@cs.technion.ac.il

Abstract

The goal of the alignment problem *is to align a (given) point cloud $P = \{p_1, \cdots, p_n\}$ to another (observed) point cloud $Q = \{q_1, \cdots, q_n\}$. That is, to compute a rotation matrix $R \in \mathbb{R}^{3\times3}$ and a translation vector $t \in \mathbb{R}^3$ that minimize the sum of paired distances between every transformed point $Rp_i - t$, to its corresponding point q_i, over every $i \in \{1, \cdots, n\}$. A harder version is the* registration problem*, where the correspondence is unknown, and the minimum is also over all possible correspondence functions from P to Q. Algorithms such as the Iterative Closest Point (ICP) and its variants were suggested for these problems, but none yield a provable non-trivial approximation for the global optimum.*

We prove that there always *exists a "witness" set of* 3 *pairs in $P \times Q$ that, via novel alignment algorithm, defines a constant factor approximation (in the worst case) to this global optimum. We then provide algorithms that recover this witness set and yield the first provable constant factor approximation for the: (i) alignment problem in $O(n)$ expected time, and (ii) registration problem in polynomial time. Such small witness sets exist for many variants including points in d-dimensional space, outlier-resistant cost functions, and different correspondence types.*

Extensive experimental results on real and synthetic datasets show that, in practice, our approximation constants are close to 1 *and our error is up to x*10 *times smaller than state-of-the-art algorithms.*

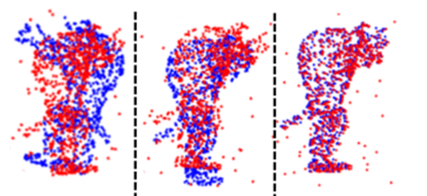

Figure 1: Registration visualization using the Armadillo model with $n = 1000$ points and $k = 20\%$ outliers. (Left) $\texttt{ICP}(P, Q)$, (middle) $\texttt{P-ICP}(P, Q, \text{cost}, \gamma)$, (right) $\texttt{P-ICP-Refined}(P, Q, \gamma, \text{cost})$. cost is the SSD with a threshold M-estimator and $\gamma = 3000$; see Section 3.2.

1. Introduction

Consider the set P of known 3D landmarks mounted on a car, and the set Q of the same 3D landmarks as currently observed via an external 3D camera, say, a few seconds later. Suppose that we wish to compute the new car's position and orientation, relative to its starting point. These can be deduced by recovering the rigid transformation (rotation and translation) that align P to Q. In this *alignment problem*, we assume that the correspondence (matching) between every point in P to Q is known. When this matching is unknown, and needs to be computed, the problem is known as the *registration problem*. It is a fundamental problem in computer vision [33, 28, 40, 47] with many applications in robotics [27, 37, 13] and autonomous driving [51].

Alignment. In the alignment problem the input consists of two ordered sets $P = \{p_1, \cdots, p_n\}$ and $Q = \{q_1, \cdots, q_n\}$ in $\mathbb{R}^d$, where $d = 3$ in the previous application, and the goal is to minimize

$$\sum_{i=1}^{n} D(Rp_i - t, q_i), \tag{1}$$

over every *alignment* (rigid transformation) (R, t) consisting of a rotation matrix $R \in \mathbb{R}^{d\times d}$ (an orthogonal matrix whose determinant is 1), and a translation vector $t \in \mathbb{R}^d$, and where $D(p, q) = \|p - q\|$ is the Euclidean (ℓ_2) distance between a pair of points $p, q \in \mathbb{R}^d$. Here, the sum is over the distance between every point $p_i \in P$ to its corresponding point $q_i \in Q$. This correspondence may be obtained using some auxiliary information, like point-wise descriptors e.g., SIFT [25], visual tracking of points [36, 44], or the use of predefined shapes and features [32, 39].

To our knowledge, the only provable approximation to the optimal *global minimum* of (1) is for its variant where $D(p, q)$ is replaced by $\ell(D(p, q)) = \|p - q\|^2$, i.e., *squared* Euclidean distance. In this special case, the optimal solution

*Table 1: **Example contributions.** Variants of the problems (1)–(3) that we approximate in this paper, either using: (i) Theorem 3 (known correspondence), or (ii) Theorem 5 (unknown correspondence). Let $P = \{p_1, \cdots, p_n\}$ and $Q = \{q_1, \cdots, q_n\}$ be two sets of points in $\mathbb{R}^d$, let $z, r, T > 0$, and let $w = d^{\left|\frac{1}{z}-\frac{1}{2}\right|}$. Formally, we wish to minimize $\text{cost}(P, Q, (R, t)) = f(\ell(D(Rp_1 - t, q_1)), \cdots, \ell(D(Rp_n - t, q_n)))$ for functions $D : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty)$, $\ell : [0, \infty) \to [0, \infty)$ and $f : \mathbb{R}^n \to [0, \infty)$ as in Definition 2. Rows marked with a $\star$ can also be approximated in linear time with high probability and bigger approximation factors, using Theorem 4.*

Use case	$f(v)$	$\ell(x)$	$D(p,q)$	Optimization Problem $\text{cost}(P, Q, (R, t))$	Approximation Factor	Matching m Necessary as input?
Sum of distances $\star$	$\\|v\\|_1$	x	$\\|p-q\\|_2$	$\sum_{i=1}^{n} \\|Rp_i - t - q_{m(i)}\\|$	$(1+\sqrt{2})^d$	No
Sum of squared distances $\star$	$\\|v\\|_1$	x^2	$\\|p-q\\|_2$	$\sum_{i=1}^{n} \\|Rp_i - t - q_{m(i)}\\|^2$	$(1+\sqrt{2})^{2d}$	No
Sum of distances with noisy data using M-estimators	$\\|v\\|_1$	$\min\{x, T\}$	$\\|p-q\\|_2$	$\sum_{i=1}^{n} \min\{\\|Rp_i - t - q_{m(i)}\\|, T\}$	$(1+\sqrt{2})^d$	No
Sum of ℓ_z distances to the power of r $\star$	$\\|v\\|_1$	x^r	$\\|p-q\\|_z$	$\sum_{i=1}^{n} \\|Rp_i - t - q_{m(i)}\\|_z^r$	$w^r(1+\sqrt{2})^{dr}$	No
Sum of ℓ_z distances to the power of r with $k \geq 1$ outliers	Sum of the $n-k$ smallest entries of v	x^r	$\\|p-q\\|_z$	$\sum_{i \in S \subset \{1,\cdots,n\}, \|S\|=n-k} \\|Rp_i - t - q_{m(i)}\\|_z^r$	$w^r(1+\sqrt{2})^{dr}$	Yes

is unique and easy to compute: t is simply the vector connecting the two centers of mass of P and Q, and $R \in \mathbb{R}^{d \times d}$ can be computed using Singular Value Decomposition [14] as described in [22]. There has been a long line of work to handle this problem also in the presence of outliers; see e.g., [6, 52, 49]. Many of those works are RANSAC-type algorithms [12]. This paper gives the first provable non-trivial approximation algorithm for (1), while also handling an even wider range of functions.

Registration. The registration problem does not assume the correspondence between P and Q is given, that is, we do not know which point in Q matches $p_i \in P$. Therefore, besides the rigid motion, the correspondence needs also to be extracted based solely on the two given point clouds, resulting in a much more complex problem with a large number of local minima; see Fig. 1. Formally, it aims to minimize

$$\sum_{i=1}^{n} \ell\left(D(Rp_i - t, q_{m(i)})\right), \qquad (2)$$

over every alignment (R, t) and correspondence function $m : \{1, \ldots, n\} \to \{1, \ldots, n\}$; see recent survey [41]. Here, a natural selection for ℓ is $\ell(x) = x^2$. The set Q here is assumed to be of size n for simplicity only, but can be of any different size.

Unlike (1), we do not know a provable approximation to (2), even for $\ell(x) = x^2$. The most commonly used solution for this problem, both in academy and industry, is the Iterative Closest Point (ICP) heuristic [4]. Our main contribution is a provable alternative to the ICP which approximates the global optimum of this harder problem.

More complex cost functions. When dealing with real-world data, noise and outliers are inevitable. One may thus consider alternative cost functions, rather than the sum of squared distances (SSD) above, due to its sensitivity to such corrupted input. A natural more general cost function would be to pick e.g., $\ell(x) = x^r$ for $r > 0$, which is more robust to noise when $r \in (0, 1]$. Alternatively, for handling outliers, a more suitable function would be $\ell(x) = \min\{x, T\}$ for some threshold $T > 0$, or the common Tuckey or Huber losses, or any other robust statistics function [17].

To completely ignore these (unknown) faulty subsets of some paired data we may consider solving

$$\min_{(R,t)} \sum_{i \in S \subset \{1,\cdots,n\}, |S|=n-k} \ell\left(D(Rp_i - t, q_{m(i)})\right), \qquad (3)$$

where $k \leq n$ is the number of outliers to ignore.

In this paper we suggest a general framework for provably approximating the global minimum of the alignment and registration problems, including formulations (1)–(3).

1.1. Related Work

The most common method for solving the registration problem in (2), for $\ell(x) = x^2$, is the ICP algorithm [7, 4]. The ICP is a local optimization technique, which alternates, until convergence, between solving the correspondence problem and the rigid alignment problem. Over the years, many variants of the ICP algorithm have been suggested; see survey in [38] and references therein. However, these methods usually converge to local and not global minimum if not initialized properly.

Estimation maximization approaches. To overcome the ICP limitations, probabilistic methods have been suggested, making use of GMMs, treating one point set as the GMM centroids, and the other as data points [18, 46, 8, 26, 29, 16, 5]. This category also includes the widely used Coherent Point Drift (CPD) method [31].

Learning-based approaches. Learning dedicated features for this task was shown to enhance the output alignment [45]. In [3], a deep learning model was combined with a modified version of the known Lukas & Kanade algorithm. Recently, an unsupervised deep learning based approach was proposed in [15].

Geometric and alternative approaches. Some works, e.g., [1, 34, 2], utilize techniques from computational geometry to devise a solution. [1, 34] also provide provable guarantees. Other results use a Branch and Bound scheme to compute the global minimum [35, 10, 50]. The work [30] tackles the problem using a smart indexing data organization. Some results use the Fourier domain [28], and use correlation of kernel density estimates (KDE) [42]. However, the above methods scale poorly as the input size increases.

Common limitations. The previously mentioned methods share similar properties and either (i) support only the simple sum of squared distances function or $d = 3$, (ii) they converge to a local minima due to bad initialization, (iii) give optimality guarantees, if any, only on a sub-task of the registration pipeline, and lack such guarantees relative to the global optimum of the registration problem, (iv) their convergence time is impractical or depends on the data itself, or (v) require a lot of training data. To our knowledge, no provable approximation algorithms have been suggested for tackling (2), even for $\ell(x) = x$.

Coresets. Some works suggest compressing the input point clouds into a small subset with a provable bound on the compression error. Existing and inefficient algorithms can thus run much faster. An error-less compression was suggested in [32] for the alignment problem; see more examples in [20]. While our method draw inspiration from the approximation techniques used in developing coresets, our paper suggests a "witness set", and not a coreset.

1.2. Our Contribution

(i) A novel alignment algorithm that given a specific set of d points from P and corresponding d points from Q, which we call a witness set, yields a provable constant factor approximation. We also prove that *every* input pair of point clouds in $\mathbb{R}^d$ admits such a witness set for all the versions of the alignment and registration problem, including problems (1)–(3); see Theorem 1.

(ii) RANSAC-style algorithms for recovering such witness sets for both the alignment and registration problems and their variants e.g., (1)–(3). Extensive experimental results on synthetic and real-world datasets demonstrate the effectiveness and accuracy of our suggested algorithms, as compared to state of the art methods; see Section 3. The results show that the approximation factor obtained in practice is much smaller than the theoretically predicted factor. We provide full open-source code for our algorithms [21].

(iii) A formal proof that running our suggested algorithms for a polynomial number of iterations yields a provable constant factor approximation for the alignment and registration problems, including sum of distances to the power of $r > 0$ and sum of M-estimators; see Theorems 3 and 5, Table 1, and Algorithms 2 and 4.

(iv) A probabilistic *linear time* algorithm for solving the alignment problem, which also supports e.g., sum of M-estimators; see Algorithm 3 and Theorem 4.

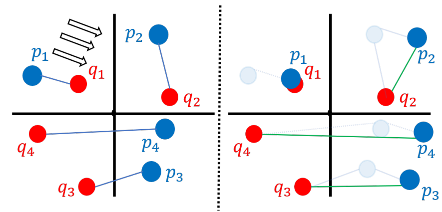

Figure 2: (Left): Two corresponding sets of points P (in blue) and Q (in red), where p_1 and q_1 have the smallest distance among all pairs. (Right): Translating P by $t = p_1 - q_1$ (i.e., p_1 now intersects q_1). By the triangle inequality, each distance $\|p_i - t - q_i\|$ (green lines) is at most $2 \cdot \|p_i - q_i\|$ (blue lines).

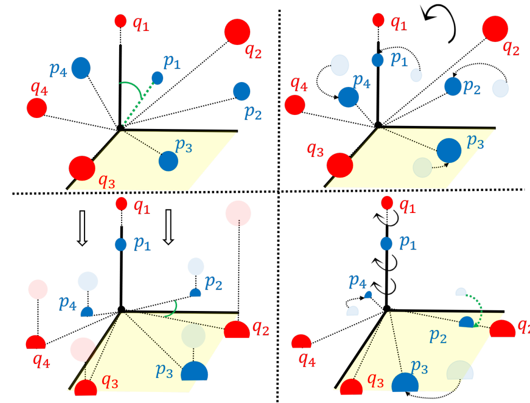

Figure 3: (Top left): Two sets of corresponding points P (in blue) and Q (in red). (Top right): Rotating P such that some $p_1 \in P$ aligns with its corresponding $q_1 \in Q$. (Bottom left): Projecting the rotated set P and the set Q onto the plane orthogonal to q_1. (Bottom right): Rotating the projected P such that one of its points aligns with its corresponding point from Q. Observe that the initial aligned pair of points (p_1, q_1) are not affected by the proceeding steps.

1.3. Novel Technique: Witness Set

We now introduce our novel technique. We first assume the correspondence between P and Q is given. We then generalize to the case with unknown correspondence.

Our main technical result is that for every *corresponding* ordered point sets $P = \{p_1, \cdots, p_n\}, Q = \{q_1, \cdots, q_n\} \subseteq \mathbb{R}^d$, every cost function cost which satisfies some set of properties (see Definition 2), and every possible alignment (R^*, t^*), there is a subset of P and a subset of Q, both of size equal to the dimension d, which we call a *witness set*. Using those subsets, our algorithm can determine an alignment (R', t'), that approximates the cost of (R^*, t^*), i.e., $\text{cost}(P, Q, (R', t')) \leq c \cdot \text{cost}(P, Q, (R^*, t^*))$, for small constant $c > 0$. Here, cost assigns a non-negative value for every pair of input point sets and alignment, and (R^*, t^*) can be the globally optimal (unknown) alignment.

For the sake of analysis only, we assume that (R^*, t^*) is known beforehand. The proof assumes an initial position of the point clouds where (R^*, t^*) has already been applied to P. It then applies a series of steps which alter this initial alignment (R^*, t^*), until a different alignment (R', t') is obtained, where a (witness) set of points from P and Q satisfies a sufficient number of known constraints, making it feasible (given this witness set) to recover (R', t'). Each step in this series is guaranteed to approximate the cost of its preceding step. Hence, the cost of (R', t') approximates the initial (optimal) cost of (R^*, t^*). The steps are as follows:

(i) Consider the set P' obtained by applying the optimal (unknown) alignment (R^*, t^*) to P. Now, consider the single corresponding pair of points $p' = R^*p - t^* \in P'$ and $q \in Q$ which have the closest distance $\|p' - q\|$ between them among all matched pairs. Using the triangle inequality, one can show that translating the set P' by $p' - q$ (that is, such that p' now intersects q) would not increase the pairwise distances of the other pairs of points by more than a multiplicative factor of 2; see Fig. 2. Hence, we proved the existence of a translation t' of P', where some $p \in P$ intersects its corresponding $q \in Q$, and where the cost is larger than the initial optimal cost by at most a constant factor. Now, assume that p' and q are located at the origin.

(ii) Similarly, we prove there is a corresponding pair of points $p' \in P', q \in Q$ such that aligning their direction vectors via a rotation R', i.e., $R' \frac{p'}{\|p'\|} = \frac{q}{\|q\|}$, would increase the pairwise distances of the other pairs by at most a small factor. p' and q are the pair with the smallest angle between them.

(iii) We can repeat step (ii) above iteratively as follows: Find such a pair (p', q), align their direction vectors, project the two sets of points onto the hyperplane orthogonal to the direction vector of q, and repeat at most $d - 2$ times. Such a projection insures that the next uncovered rotation will maintain the alignment of (p', q); see Fig. 3. Each such step proves *the existence* of yet another corresponding pair of points which contribute at least 1 constraint on R', without damaging the cost by more than a constant factor. Hence, there exist $d - 1$ pairs of points which uniquely determine our approximated rotation R'. We call the d (unknown) pairs from the steps above a *witness set*. Given a witness set, the approximated alignment (R', t') can be recovered.

Recovering such a witness set requires recovering d points from P (where d is the dimension of P), which, using the known correspondence, uncover the d corresponding points from Q. When the correspondence is unknown, the minor difference is that we need to recover d points from P as well as d independent points from Q; see Theorem 5.

2. Provable Approximations

We now prove the existence of a witness set for the alignment and registration problems and their variants presented above. We then present algorithms that recover such a witness set for each of the problems. Due to lack of space, all our proofs are placed in the supplementary material.

Notation. We denote $[n] = \{1, \cdots, n\}$ for any integer $n \geq 1$. We assume every vector is a column vector. Let $\text{SO}(d)$ be the set of all rotation matrices in $\mathbb{R}^d$. For $t \in \mathbb{R}^d$ and $R \in \text{SO}(d)$, the pair (R, t) is called an *alignment*. We define $\text{ALIGNMENTS}(d)$ to be the union of all possible d-dimensional alignments. A *correspondence* (or matching) function is simply a function $m : [n] \to [n]$. For a correspondence function m and an ordered set $P = \{p_1, \cdots, p_n\}$, we define $P_{[m]} = \left\{p_{m(1)}, \cdots, p_{m(n)}\right\}$ as a new ordered set obtained from P after reordering its elements according to m.

2.1. Existence of a Witness Set

In what follows we present our main alignment algorithm and our main technical result; see Algorithm 1 and Theorem 1. Theorem 1 proves the existence of some witness set that, when plugged into Algorithm 1, produces an alignment with some provable guarantees.

Overview of Algorithm 1. Algorithm 1 gets as input two sets $\{p_1, \cdots, p_d\}$ and $\{q_1, \cdots, q_d\}$ in $\mathbb{R}^d$, and implements the scheme described in Section 1.3. At Line 1, we translate both sets so that p_d and q_d intersect at the origin. At Line 4 we compute a rotation matrix S that aligns the directions of p_1 and q_1. In Lines 5– 6 we compute an orthogonal matrix W whose column space spans the $(d - 1)$-dimensional subspace π orthogonal to q_1 and project $Sp_2, \cdots, Sp_d, q_2, \cdots, q_d$ onto π, and repeat $d - 1$ times. Hence, at the i'th iteration, we compute a rotation S that aligns the directions of p_i (after $i - 1$ rotations and projections) and q_i (after $i - 1$ projections), but also maintains the alignment of the previously aligned pairs $(p_1, q_1), \cdots, (p_{i-1}, q_{i-1})$. Such a matrix exists since $p_i, \cdots, p_d, q_i, \cdots, q_d$ are, by construction, orthogonal to $p_1, \cdots, p_{i-1}, q_1, \cdots, q_{i-1}$; see Fig. 3. We output an alignment which replicates the composition of the steps above.

Overview of Theorem 1. Consider $P = \{p_1, \cdots, p_n\}$ and $Q = \{q_1, \cdots, q_n\}$ and any variant of either the alignment or registration problems, e.g., (1)–(3). Now, assume that (R^*, t^*) and m^* are respectively the globally optimal alignment and correspondence function for the task at hand. Observe that, in the alignment problem, m^* is given as input. Theorem 1 proves the existence of a set of d points

Algorithm 1: ALIGN$(\{p_1, \cdots, p_d\}, \{q_1, \cdots, q_d\})$

Input : Two sets of points that each spans a $d-1$ dimensional subspace in $\mathbb{R}^d$.
Output: An alignment (R, t); see Theorem 1

```
p_i := p_i − p_d and q_i := q_i − q_d for every i ∈ [d]
R := the identity matrix in ℝ^d
for every z ∈ [d − 1] do
    S := an arbitrary rotation matrix that satisfies
      S p_z / ‖p_z‖ = q_z / ‖q_z‖, and S p_i = p_i for every i ∈ [z − 1].
    W := an arbitrary matrix in ℝ^{d×(d−1)} such that
      [W | q_z/‖q_z‖] ∈ ℝ^{d×d} forms a basis of ℝ^d.
    p_i := W W^T S p_i and q_i := W W^T q_i, for every
      i ∈ [d] \ [z]
    p_z := S p_z
    R := S R
t := R p_d − q_d
return (R, t)
```

from P and d points from Q that, when plugged into Algorithm 1, produce an alignment (R, t) which guarantees:

$$\left\|Rp_i - t - q_{m^*(i)}\right\| \leq \sigma \cdot \left\|R^*p_i - t^* - q_{m^*(i)}\right\|, \quad \forall i \in [n]$$

for some small $\sigma \geq 1$. For $d = 3$ the constant is $\sigma < 15$. In other words, using (R, t) we can approximate each of the n pairwise distances of the optimal alignment (R^*, t^*). In the registration problem, m^* can be recovered afterwards, e.g., via nearest neighbour algorithm. Theorem 1 thus successfully decouples the two problems of recovering the alignment and recovering the correspondence function.

Theorem 1 (Witness sets)**.** *Let $P = \{p_1, \cdots, p_n\}$ and $Q = \{q_1, \cdots, q_n\}$ be two ordered sets each of n points in $\mathbb{R}^d$. Then, for every alignment (R^*, t^*) and matching function m^*, there exist $P' \subseteq P$ and $Q' \subseteq Q$ of size $|P'| = |Q'| = d$ such that the output (R, t) of the call* ALIGN(P', Q') *to Algorithm 1 satisfies the following for every $i \in [n]$:*

$$\left\|Rp_i - t - q_{m^*(i)}\right\| \leq (1 + \sqrt{2})^d \cdot \left\|R^*p_i - t^* - q_{m^*(i)}\right\|$$

Furthermore, (R, t) is computed in $O(d^3)$ time.

In Section 2.2 we prove that individually approximating the pairwise distances, as in Theorem 1, implies an immediate approximation to a wide range of cost functions.

While Theorem 1 guarantees the existence of at least one such witness set, empirically we have observed that many subsets of P and Q serve as good witness sets, in the sense that they produce approximation factors smaller than predicted in the theorem. Those factors are usually even close to 1. Hence, in Sections 2.3– 2.4 we apply RANSAC-type algorithms to recover a witness set. Theorem 1 also implies that running the suggested algorithms for a sufficient number of iterations produces a guaranteed constant factor approximation to the global optimum of the problems at hand.

2.2. Generalization

In what follows we define a wide family of cost functions which this work tackles, including the cost functions in (1)–(3). We then show that the approximation guarantees obtained in Theorem 1 suffice to approximate each such cost function; See Table 1 for examples. In what follows, for $r > 0$, an r-log-Lipschitz function is a function that, in every dimension individually, may be large but cannot increase too rapidly (in a rate that depends on r); see formal definition in Section B at the appendix.

Definition 2 (Cost function)**.** *Let $P = \{p_1, \cdots, p_n\} \subseteq \mathbb{R}^d$ and $Q = \{q_1, \cdots, q_n\} \subseteq \mathbb{R}^d$. Let $D : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty)$ be a function that assigns a non-negative weight for each pair of points in $\mathbb{R}^d$, $\ell : [0, \infty) \to [0, \infty)$ be an r-log-Lipschitz function and $f : [0, \infty)^n \to [0, \infty)$ be an s-log-Lipschitz function. Let (R, t) be an alignment. We define*

$$\begin{aligned}&\text{cost}(P, Q, (R, t)) \\ &\quad = f\left(\ell\left(D\left(Rp_1 - t, q_1\right)\right), \cdots, \ell\left(D\left(Rp_n - t, q_n\right)\right)\right).\end{aligned}$$

From pairwise distances to complex cost functions. Observation 5 in [19] (see Section B in the appendix) states that in order to approximate a given cost function from Definition 2, relative to the globally optimal alignment and correspondence (R^*, t^*) and m^* respectively, it is sufficient to approximate, simultaneously, each of the pairwise distances $\left\|R^*p_i - t^* - q_{m^*(i)}\right\|$ for every $i \in [n]$. By Theorem 1, there exists a witness set which provides the desired approximation for the above pairwise distances. Combining the above yields a provable approximation for any cost function from Definition 2. It is only left to recover a designated witness set, which is the goal of Sections 2.3 and 2.4.

2.3. Approximations for the Alignment Problem

We now provide an algorithm which computes a provable approximation for the alignment problem (1) and its variants, i.e., when the matching between P and Q *is given*. This is by recovering a designated witness set.

Overview of Algorithm 2. Using Theorem 1 and assuming the matching between P and Q is given, one can construct a RANSAC-type algorithm that iterates over γ subsets of P and Q of size d, applies Algorithm 1 to each two such corresponding subsets to obtain a candidate alignment, and returns the alignment that minimizes the cost function at hand. Algorithm 2 implements the scheme above.

Theorem 3. *Let P and Q be two ordered sets of n points in $\mathbb{R}^d$, $\gamma \in \Omega(n^d)$, $z > 0$, and $w = d^{\left|\frac{1}{z} - \frac{1}{2}\right|}$. Let* cost*, r and s be as defined in Definition 2 for $D(p, q) =$*

Algorithm 2: Approx-Alignment$(P, Q, \gamma, \text{cost})$

Input : A pair of sets $P \subseteq \mathbb{R}^d$ and $Q \subseteq \mathbb{R}^d$, number of iterations $\gamma > 0$, and a cost function.

Output: An alignment (R, t); see Theorem 3

$M := \emptyset$.

$I :=$ randomly sample, with no repetition, a set of γ tuples of d distinct indices from $[n]$. // $|I| = \gamma$

for *every* $(i_1, \cdots, i_d) \in I$ **do**

  $(R', t') :=$ Align$(\{p_{i_1}, \cdots, p_{i_d}\}, \{q_{i_1}, \cdots, q_{i_d}\})$ // see Algorithm 1

  $M := M \cup \{(R', t')\}$

$(R, t) \in \arg\min_{(R', t') \in M} \text{cost}(P, Q, (R', t'))$

return (R, t)

$\|p - q\|_z$. *Let (R, t) be the output of a call to* Approx-Alignment$(P, Q, \gamma, \text{cost})$*; See Algorithm 2. Then,*

$$\text{cost}(P, Q, (R, t)) \leq w^{rs} \cdot (1 + \sqrt{2})^{drs} \cdot \min_{(R', t')} \text{cost}(P, Q, (R', t')),$$

where the minimum is over every $(R', t') \in$ Alignments*. Moreover, (R, t) is computed in $n^{O(d)}$ time.*

2.3.1 Run-Time Improvement

We now propose a randomized algorithm (see Algorithm 3) with the same goal as Algorithm 1, that succeeds with constant probability. By running this algorithm for a constant number of times, we can recover an alignment that, with probability approaching 1, has the same guarantees as the output of Algorithm 2. However, this new randomized algorithm requires linear, rather than polynomial, time.

Overview of Algorithm 3. Unlike Algorithm 1, which expects to receive a witness set as input, Algorithm 3 takes as input two full point clouds, and internally identifies a potential witness set. Intuitively, points in P with larger norm negatively affect our cost function more than points of smaller norm, when misaligned properly; see Fig. 11 at the appendix. Algorithm 3 thus samples a pair of corresponding points (p, q) with probability that depends on the norm of p and rotates P to align the direction vectors of p and q. Then, similarly to Algorithm 1, it projects the sets onto the hyperplane orthogonal to q and repeats.

Theorem 4. *Let P and Q be two ordered sets of n points in $\mathbb{R}^d$ and $z > 0$. Let* cost *be as in Definition 2 for $f = \|v\|_1$, some r-log Lipschitz function ℓ and $D(p, q) = \|p - q\|_z$. Let (R, t) be an output of a call to* Prob-Align(P, Q, r)*; see Algorithm 3. Then, with probability at least $\frac{1}{2^d}$,*

$$\text{cost}(P, Q, (R, t)) \leq \sigma \cdot \min_{(R', t') \in \text{Alignments}(d)} \text{cost}(P, Q, (R', t')),$$

Algorithm 3: Prob-Align(P, Q, r)

Input : A pair of sets $P = \{p_1, \cdots, p_n\}$ and $Q = \{q_1, \cdots, q_n\}$ in $\mathbb{R}^d$ and $r > 0$.

Output: A rotation matrix; see Theorem 4

Sample an index $k \in [n]$ uniformly at random

$p := p - p_k$ for every $p \in P$

$q := q - q_k$ for every $q \in Q$

$J := \{k\}$ and $R :=$ the d-dimensional identity matrix

for *every* $z \in [d - 1]$ **do**

  $w_i := \frac{\|p_i\|^r}{\sum_{j \in [n]} \|p_j\|^r}$ for every $i \in [n]$.

  Randomly sample an index $j \in [n] \setminus J$, where $j = i$ with probability w_i.

  $S :=$ an arbitrary rotation matrix that satisfies $\frac{Sp_j}{\|p_j\|} = \frac{q_j}{\|q_j\|}$ and $Sp_i = p_i$ for every $i \in J$.

  $W :=$ a matrix in $\mathbb{R}^{d \times (d-1)}$ such that $\left[W \mid \frac{q_j}{\|q_j\|}\right] \in \mathbb{R}^{d \times d}$ forms a basis of $\mathbb{R}^d$.

  $J := J \cup \{j\}$

  $p := Sp$ for every $p \in P$

  $p := WW^T p$ for every $p \in P \setminus \{p_i \mid i \in J\}$

  $q := WW^T q$ for every $q \in Q \setminus \{q_i \mid i \in J\}$

  $R := SR$

$t := Rp_k - q_k$

return (R, t)

for a constant σ that depends on d and r. Furthermore, (R, t) is computed in $O(nd^3)$ time.

Success probability. For the usual case of $d = 3$, the success probability of Algorithm 3 is at least $1/8$. Hence, repeating Algorithm 3 for less than 6 repetitions amplifies the success probability in Theorem 4 to more than $1/2$.

2.4. Approximations for the Registration Problem

As explained in Section 1.3, a witness set from P and Q also exists in the much harder variant where the matching between P and Q is unknown. We now provide an algorithm that, for any given variant of the registration problem, can recover a witness set. The formal statement is given in Theorem 5, which is one of our main contributions.

Overview of Algorithm 4. Unlike Algorithm 2 which samples d indices used to index both P and a Q, we now have to independently sample d indices for points in P as well as d indices for points in Q. Furthermore, we need to compute the nearest neighbour matching for every candidate alignment returned by Algorithm 1, before evaluating the cost function. Algorithm 4 applies the above scheme γ times and returns the alignment and matching function that minimize the given cost function.

Theorem 5. *Let $P = \{p_1, \cdots, p_n\}$, $Q = \{q_1, \cdots, q_n\}$ be two ordered sets of n points in $\mathbb{R}^d$, $\gamma \in \Omega(n^{2d})$, $z > 0$,*

Algorithm 4: ALIGN-AND-MATCH$(P, Q, \gamma, \text{cost})$

Input : A pair of sets $P \subseteq \mathbb{R}^d$ and $Q \subseteq \mathbb{R}^d$, number of iterations $\gamma > 0$, and a cost function.

Output: An alignment and a matching function; see Theorem 5

$M := \emptyset$.

$I :=$ randomly sample, with no repetition, a set of γ tuples of $2d$ indices $(i_1, \cdots, i_d, j_1, \cdots, j_d)$ from $[n]$, such that $i_1, \cdots, i_d$ are distinct and $j_1, \cdots, j_d$ are distinct. // $|I| = \gamma$

for *every* $(i_1, \cdots, i_d, j_1, \cdots, j_d) \in I$ **do**

  $(R', t') :=$ ALIGN$(\{p_{i_1}, \cdots, p_{i_d}\}, \{q_{j_1}, \cdots, q_{j_d}\})$ // see Algorithm 1

  $M := M \cup \{(R', t', \text{NN}(P, Q, (R', t')))\}$ /* $\text{NN}(P, Q, (R, t))$ is the nearest neighbour matching between Q, and P after applying (R, t). */

$(\tilde{R}, \tilde{t}, \tilde{m}) \in \underset{(R', t', m') \in M}{\arg\min} \ \text{cost}\left(P_{[m']}, Q, (R', t')\right)$.

return $(\tilde{R}, \tilde{t}, \tilde{m})$

and $w = d^{\left|\frac{1}{z} - \frac{1}{2}\right|}$. *Let* cost *and* r *be as in Definition 2 for* $D = \|p - q\|_z$ *and* $f(v) = \|v\|_1$. *Let* $(\tilde{R}, \tilde{t}, \tilde{m})$ *be the output of a call to* ALIGN-AND-MATCH$(P, Q, \gamma, \text{cost})$*; See Algorithm 4. Then, for* $c = w^r (1 + \sqrt{2})^{dr}$*, we have*

$$\text{cost}\left(P_{[\tilde{m}]}, Q, (\tilde{R}, \tilde{t})\right) \leq c \cdot \min_{(R,t,m)} \text{cost}\left(P_{[m]}, Q, (R, t)\right),$$

where the minimum is over every alignment (R, t) *and permutation* m*. Moreover,* $(\tilde{R}, \tilde{t}, \tilde{m})$ *is computed in* $n^{O(d)}$ *time.*

Substituting $z = r = 2$ in Theorem 5 yields a provable approximated alternative to ICP.

Comparison to RANSAC. RANSAC has some similarity to Algorithms 2 and 4 above in the sense that they both randomly sample points from P and Q. However, while RANSAC recovers a candidate alignment via common least squares on such candidate d pairs (see Section 1), our algorithms utilize a novel method (Algorithm 1) which is not necessarily optimal for those d pairs, but will be (almost) globally optimal for the overall cost of all the n pairs. In theory, unlike our algorithms, RANSAC does not guarantee global optimality; see comparison in Section 3.1.

Our approximation constants. While the approximation constants in Theorems 3- 5 above might seem large, they are only roughly < 14 in the pessimistic worst-case theory, they are smaller than 2 in practice, and can be obtained much faster than the suggested time; see Section 3.

3. Experimental Results

We now apply our algorithms to solve either the alignment or the registration problems. Additional experiments on real-world scans from the SUN3D dataset [48] are placed in Section F at the appendix.

Datasets. We used the Bunny, Armadillo, and Asian Dragon models from the Stanford 3D scanning repository [9, 23, 43]. Those models were scaled to $[-0.5, 0.5]^3$ due to the constraints of some competing methods (e.g., `GO-ICP`). We also used a synthetic dataset comprising uniformly sampled d-dimensional points in $[-0.5, 0.5]^d$.

Generating P and Q. In all experiments, given some data model (real or synthetic), we uniformly sample n points named Q. An alignment (R, t) is generated, where t is uniformly sampled such that $\|t\| \leq 0.1$ and R rotates the data around each axis by an angle uniformly sampled from $[-\pi, \pi]$. P is then obtained by applying (R, t) to Q and adding Gaussian noise with zero mean and σ^2 variance. In Section 3.2 we also apply a random shuffle to P.

Evaluation. We present two evaluation metrics: (i) The value of the minimized cost function itself, e.g., the value of (1) or (3). If not given, the optimal correspondence is trivially computed, after applying (R, t) to P, via nearest neighbor. (ii) Rotation and translation errors: $\left\|R^T R^* - I\right\|_F$ and $\|t - t^*\|_2$, where (R^*, t^*) and (R, t) are the ground truth and the recovered alignments, respectively. Every experiment in this section was conducted 20 times and averaged. The variance is presented in the graphs.

3.1. Alignment Experiments

Here, we assume the correspondences function is given. We applied three algorithms: (i) `P-RANSAC`: Provable RANSAC - an implementation in Python of Algorithm 2, (ii) `RANSAC`: A RANSAC scheme equipped with the common least squares solution to recover an alignment, and (iii) `TEASER++`: The state of the art TEASER++ [49].

In Fig. 4 we compare the algorithms above, and in Fig. 5 we demonstrate `P-RANSAC`'s fast recovery of an alignment with an approximation constant close to 1, even in high dimensions. In the latter test, the ground truth solution (R^*, t^*) is computed via SVD as explained in Section 1.1.

Discussion. Fig. 4 demonstrates that Algorithm 2 outperforms state of the art methods, in multiple common metrics. Fig. 5 shows that it suffices for our algorithm, in practice, to sample roughly 40 subsets (in $d = 3$) until an error of at most x1.5 the globally minimal cost is obtained. Furthermore, recall that Algorithm 3 provides a probabilistic alternative for Algorithm 2, by reducing number of iterations in the cost of larger approximation constants. However, Algorithm 2 was sufficient in practice as it produced very small approximation constants in very few iterations.

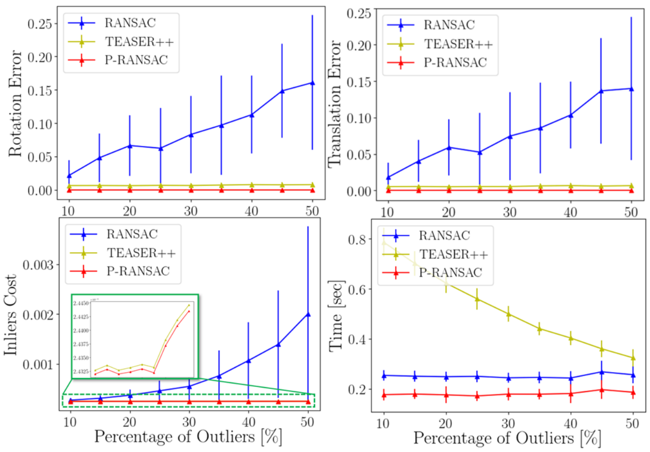

Figure 4: The Bunny model, $n = 2000$ points and $\sigma^2 = 0.009$ noise variance were used. The cost in our algorithm was SSD with M-estimator $\min\{\|p-q\|^2, 1\}$. The inliers cost is the mean error over the ground truth inlier pairs.

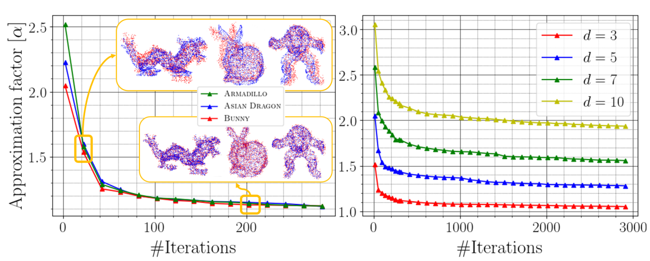

Figure 5: The approximation quality α (the cost of `P-RANSAC` *divided by the ground truth) as a function of the number of iteration γ. (Left): Bunny, Armadillo, and Asian Dragon models. The obtained alignment is also visualized. (Right): Synthetic data. In both figures $n = 2500$ and $\sigma^2 = 0.01$. The cost in our algorithm was the SSD.*

3.2. Registration Experiments

In this section we compared the following algorithms: **(i)** `ICP`(P, Q) - An implementation of the ICP algorithm [7]. **(ii)** `P-ICP`$(P, Q, \gamma, \text{cost})$ - Provable ICP; A parallelized implementation in Python of Algorithm 4. **(iii)** `P-ICP-Refined`$(P, Q, \gamma, \text{cost})$ - Applying the output alignment of `P-ICP` to the set P, then refining the alignment via single `ICP` run, on Q and the transformed P. As the ICP is guaranteed to converge to a local minimum, it can only help reduce the cost of our (approximately optimal) output result to the closest (hopefully global) minimum. **(iv)** `CPD`(P, Q) - The Coherent Point Drift algorithm [31]. **(v)** `GO-ICP`(P, Q) - The common ICP variant [50].

We tested multiple cost functions. Fig. 6 and Fig. 7 present the results for noisy input data, and data containing outliers respectively. Visual comparison is shown in Fig. 1.

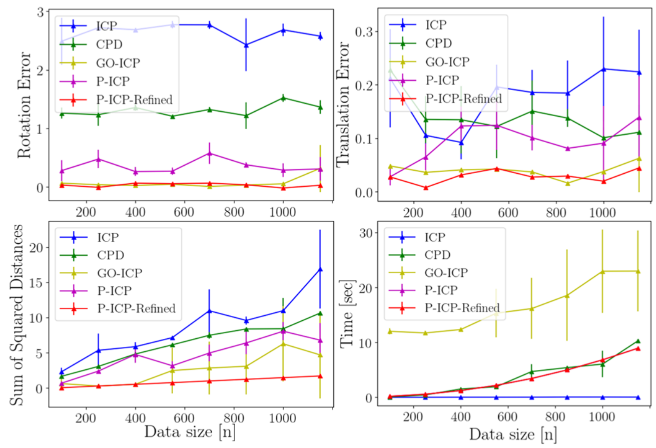

Figure 6: Armadillo model with $\sigma^2 = 0.01$ noise variance. The SSD cost function was used in our algorithms. The test was executed on the AWS platform, on a c5a.8xlarge machine with 32 CPUs.

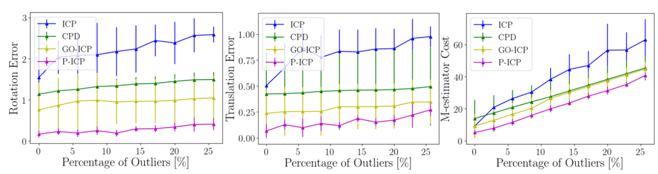

Figure 7: Robustness to outliers using the Armadillo model. $n = 800$ was used. Noise with variance $\sigma^2 = 1$ was added to k percentage of the points in P, which are considered as outliers. The SSD with M-estimator $\min\{\|p-q\|^2, 0.2\}$ was used in our algorithms. The computational time was roughly constant for each method for all tested k values, and is presented in Fig. 6 at $n = 800$.

Discussion. Fig. 6 demonstrates the accuracy of our algorithms, which yield an error smaller by x2-x10 than other methods, while also being among the fastest. Fig. 7 and 1 demonstrate our robustness to outliers in practice, due to our provable approximation to M-estimators cost functions.

4. Conclusions, Limitations, and Future Work

We present provable and practical non-trivial approximation algorithms for the alignment and registration problems and their hard variants. The algorithms rely on our proof that a witness set, which determines an approximated alignment, exists for both problems. Experiments show that our algorithms are efficient in practice, produce smaller errors, and are more stable than competing methods. The main limitation of our algorithm is their high running-time dependency on the dimension d. However, fortunately, for most applications d is a small constant. Future work includes: (i) generalizing to the non-rigid registration problem, and (ii) fast recovery of a witness set via deep learning.

References

[1] Pankaj K Agarwal and Jeff M Phillips. On bipartite matching under the rms distance. In *Proceedings of the Eighteenth Canadian Conference on Computational GeometryâĂŹ, Kingston, Canada*, pages 143–146, 2006. 3

[2] Dror Aiger, Niloy J Mitra, and Daniel Cohen-Or. 4-points congruent sets for robust pairwise surface registration. In *ACM SIGGRAPH 2008 papers*, pages 1–10. 2008. 3

[3] Yasuhiro Aoki, Hunter Goforth, Rangaprasad Arun Srivatsan, and Simon Lucey. Pointnetlk: Robust & efficient point cloud registration using pointnet. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7163–7172, 2019. 2

[4] 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 586–607. International Society for Optics and Photonics, 1992. 2

[5] 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. 2

[6] Alvaro Parra Bustos and Tat-Jun Chin. Guaranteed outlier removal for point cloud registration with correspondences. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2868–2882, 2017. 2

[7] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992. 2, 8

[8] Andrew DJ Cross and Edwin R Hancock. Graph matching with a dual-step em algorithm. *IEEE transactions on pattern analysis and machine intelligence*, 20(11):1236–1253, 1998. 2

[9] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996. 7

[10] Nadav Dym and Shahar Ziv Kovalsky. Linearly converging quasi branch and bound algorithms for global rigid registration. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1628–1636, 2019. 3

[11] Dan Feldman and Leonard J Schulman. Data reduction for weighted and outlier-resistant clustering. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 1343–1354. SIAM, 2012. 16, 17

[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. 2

[13] Miguel Angel Garcia and Agusti Solanas. 3d simultaneous localization and modeling from stereo vision. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 1, pages 847–853. IEEE, 2004. 1

[14] Gene H Golub and Christian Reinsch. Singular value decomposition and least squares solutions. In *Linear Algebra*, pages 134–151. Springer, 1971. 2

[15] Oshri Halimi, Or Litany, Emanuele Rodola, Alex M Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4370–4379, 2019. 3

[16] Jeroen Hermans, Dirk Smeets, Dirk Vandermeulen, and Paul Suetens. Robust point set registration using em-icp with information-theoretically optimal outlier handling. In *CVPR 2011*, pages 2465–2472. IEEE, 2011. 2

[17] Peter J Huber. *Robust statistics*. Springer, 2011. 2

[18] Anupam Joshi and C-H Lee. On the problem of correspondence in range data and some inelastic uses for elastic nets. *IEEE transactions on neural networks*, 6(3):716–723, 1995. 2

[19] Ibrahim Jubran and Dan Feldman. Aligning points to lines: Provable approximations. *IEEE Transactions on Knowledge and Data Engineering*, 2020. 5, 17

[20] Ibrahim Jubran, Alaa Maalouf, and Dan Feldman. Introduction to coresets: accurate coresets. *arXiv preprint arXiv:1910.08707*, 2019. 3

[21] Ibrahim Jubran, Alaa Maalouf, Ron Kimmel, and Dan Feldman. Open source code for the algorithms presented in this paper, 2021. Link for open-source code. 3

[22] Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976. 2

[23] Venkat Krishnamurthy and Marc Levoy. Fitting smooth surfaces to dense polygon meshes. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 313–324, 1996. 7

[24] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 23

[25] 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. 1

[26] Bin Luo and Edwin R Hancock. A unified framework for alignment and correspondence. *Computer Vision and Image Understanding*, 92(1):26–55, 2003. 2

[27] Martin Magnusson, Achim Lilienthal, and Tom Duckett. Scan registration for autonomous mining vehicles using 3d-ndt. *Journal of Field Robotics*, 24(10):803–827, 2007. 1

[28] Ameesh Makadia, Alexander Patterson, and Kostas Daniilidis. Fully automatic registration of 3d point clouds. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 1, pages 1297–1304. IEEE, 2006. 1, 3

[29] Graham McNeill and Sethu Vijayakumar. A probabilistic approach to robust shape matching. In *2006 International Conference on Image Processing*, pages 937–940. IEEE, 2006. 2

[30] Nicolas Mellado, Dror Aiger, and Niloy J Mitra. Super 4pcs fast global pointcloud registration via smart indexing. In *Computer Graphics Forum*, volume 33, pages 205–215. Wiley Online Library, 2014. 3

[31] Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010. 2, 8

[32] Soliman Nasser, Ibrahim Jubran, and Dan Feldman. Coresets for kinematic data: From theorems to real-time systems.

arXiv preprint arXiv:1511.09120, 2015. 1, 3

[33] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 IEEE International Symposium on Mixed and Augmented Reality*, pages 127–136. IEEE, 2011. 1

[34] Carl Olsson, Olof Enqvist, and Fredrik Kahl. A polynomial-time bound for matching and registration with outliers. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 3

[35] Carl Olsson, Fredrik Kahl, and Magnus Oskarsson. Branch-and-bound methods for euclidean registration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):783–794, 2008. 3

[36] OptiTrack. Motion capture systems. "https://www.optitrack.com". 1

[37] François Pomerleau, Francis Colas, Roland Siegwart, et al. A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends® in Robotics*, 4(1):1–104, 2015. 1

[38] François Pomerleau, Francis Colas, Roland Siegwart, and Stéphane Magnenat. Comparing icp variants on real-world data sets. *Autonomous Robots*, 34(3):133–148, 2013. 2

[39] Roman Rabinovich, Ibrahim Jubran, Aaron Wetzler, and Ron Kimmel. Cobe-coded beacons for localization, object tracking, and slam augmentation. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, pages 1367–1374. IEEE, 2020. 1

[40] Joaquim Salvi, Carles Matabosch, David Fofi, and Josep Forest. A review of recent range image registration methods with accuracy evaluation. *Image and Vision computing*, 25(5):578–596, 2007. 1

[41] Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *IEEE transactions on visualization and computer graphics*, 19(7):1199–1217, 2012. 2

[42] Yanghai Tsin and Takeo Kanade. A correlation-based approach to robust point set registration. In *European conference on computer vision*, pages 558–569. Springer, 2004. 3

[43] Greg Turk and Marc Levoy. Zippered polygon meshes from range images. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 311–318, 1994. 7

[44] Vicon. Motion systems. "https://www.vicon.com/". 1

[45] 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. 2

[46] William M Wells III. Statistical approaches to feature-based object recognition. *International Journal of Computer Vision*, 21(1-2):63–98, 1997. 2

[47] Thomas Whelan, Michael Kaess, Hordur Johannsson, Maurice Fallon, John J Leonard, and John McDonald. Real-time large-scale dense rgb-d slam with volumetric fusion. *The International Journal of Robotics Research*, 34(4-5):598–626, 2015. 1

[48] Jianxiong Xiao, Andrew Owens, and Antonio Torralba. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *Proceedings of the IEEE international conference on computer vision*, pages 1625–1632, 2013. 7, 24

[49] Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 2020. 2, 7

[50] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-icp: A globally optimal solution to 3d icp point-set registration. *IEEE transactions on pattern analysis and machine intelligence*, 38(11):2241–2254, 2015. 3, 8

[51] Ji Zhang and Sanjiv Singh. Visual-lidar odometry and mapping: Low-drift, robust, and fast. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2174–2181. IEEE, 2015. 1

[52] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *European Conference on Computer Vision*, pages 766–782. Springer, 2016. 2