

Outlier Robust ICP for Minimizing Fractional RMSD

Jeff M. Phillips
Duke University
jeffp@cs.duke.edu

Ran Liu
Duke University
ran@cs.duke.edu

Carlo Tomasi
Duke University
tomasi@cs.duke.edu

Abstract

We describe a variation of the iterative closest point (ICP) algorithm for aligning two point sets under a set of transformations. Our algorithm is superior to previous algorithms because (1) in determining the optimal alignment, it identifies and discards likely outliers in a statistically robust manner, and (2) it is guaranteed to converge to a locally optimal solution. To this end, we formalize a new distance measure, fractional root mean squared distance (FRMSD), which incorporates the fraction of inliers into the distance function. Our framework can easily incorporate most techniques and heuristics from modern registration algorithms. We experimentally validate our algorithm against previous techniques on 2 and 3 dimensional data exposed to a variety of outlier types.

1 Introduction

Aligning an input data set to a model data set is fundamental to many important problems such as scanned model reconstruction [12], structural biochemistry [19], and medical imaging [8]. The input data and the model data are typically given as a set of points. A point set may arise from laser scans of a 3D or 2D model, coordinates of atoms in a protein, positions of a lesions from a medical patient, or some other sparse representation of data. However, the relative positions of these point sets is not known, making the task of registering them nontrivial.

A popular approach to solving this problem is known as the iterative closest point (ICP) algorithm [1, 3] which alternates between finding the optimal correspondence between points, and finding the optimal transformation of one point set onto the other. As both steps reduce the distance between the point sets, this process converges, but only to a local minimum. The effectiveness, simplicity, and generality of this algorithm has led to many variations [16, 15, 4, 5, 6, 19]. For instance, the set of legal transformations can be just translations, all rigid motions, or all affine transformations. Other versions replace the

optimal correspondence between points by aligning each data point to the closest point on an implicit surface of the model data [3]. Or the traditional squared distance can be replaced with a more efficient and stable approximation to the squared distance function [11]. A slightly outdated, but excellent survey [16] evaluates many techniques.

Others have attempted to solve the global registration problem [13, 7], where for any initial alignment they attempt to find the optimal alignment between two point sets. This is often done in two steps. First find a rough global alignment by corresponding certain distinguishable feature points. Second refine the alignment with ICP.

However, all of these algorithms are vulnerable to point sets with outliers. Outliers may result from deformation of a deformable model, measurement error, spurious data that was ignored or missed in the model, partial matches because the point sets represent overlapping but not identical pieces of the same object, or interesting changes in the underlying object between time steps or among comparable objects. In short, outliers are unavoidable. Because ICP will find correspondences for all points, and then find the optimal transformation for the entire point set, the outliers will skew the alignment. Many heuristics have been suggested [5, 4] including only aligning points within a set threshold [20, 17], but most of these techniques are not guaranteed to converge, and thus can possibly go into an infinite loop, or require an expensive check to prevent this. If the fraction f of points which are outliers is known, then Trimmed ICP (TrICP) [4] can be used to find the optimal alignment of the most relevant fraction f of points. However, this fraction is rarely known a priori. If an alignment is given then RANSAC type methods [2] can be used to determine a good threshold for determining these outliers. There are also many ad hoc solutions to this problem. However, if the outliers are excluded from the data set in a particular alignment, then the alignment is no longer optimal, since those outliers which were removed influenced how the points were initially aligned.

Our contributions. Our solution to these problems is to incorporate the fraction of points which are outliers into the problem statement and into the function being optimized.

To this end, this paper makes the following contributions:

- We formalize a new distance measure between point sets which accounts for outliers: FRMSD (Section 2).
- We provide an algorithm, Fractional ICP, to optimize FRMSD (Section 3) which we prove to converge to a local optimum in the correspondence, transformation, and fraction of outliers (Section 3.2).
- We give mathematical intuition for why FRMSD aligns data points which are more likely to be inliers than outliers (Section 4 and Section 5).
- Finally, we empirically demonstrate that Fractional ICP identifies the correct alignment while simultaneously determining the outliers (Section 6).

2 Fractional RMS Distance

Consider two point sets $D, M \in \mathbb{R}^d$. The goal of this paper is to align an input data set D to a model data set M under some class of transformations, \mathcal{T} . We assume that these point sets are quite similar and there exists a strong correspondence between most points in the data. There may, however, be outliers, points in either set which are not close to any point in the other set. Let $\mu : D \rightarrow M$ match each point of D to the closest point of M .

Definition 2.1. [RMSD] *The root mean squared distance (or RMSD) is defined:*

$$\text{RMSD}(D, M, \mu) = \sqrt{\frac{1}{|D|} \sum_{p \in D} \|p - \mu(p)\|^2}$$

Problem 2.1. [minimizing RMSD] *Compute the transformation $T \in \mathcal{T}$ to minimize:*

$$\min_{T \in \mathcal{T}} \sqrt{\frac{1}{|D|} \sum_{p \in D} \|T(p) - \mu(p)\|^2}$$

Problem 2.1 is algorithmically difficult because as T varies, so does the optimal matching μ . Also, RMSD is quite susceptible to outliers because the squared distance gives a large weight to outliers. To counteract this, a specific fraction $f \in [0, 1]$ of points from D can be used in the alignment and in the distance measure between the point sets. These $f|D|$ points can be chosen to solve Problem 2.1 by selecting the points which have the smallest residual distance $r = \|p - \mu(p)\|$. Let $D_f = \{p \in D \mid |D_f| = \lfloor f|D| \text{ and } \text{RMSD}(D_f, M) \text{ is minimized}\}$. But what fraction of points should be used? We can always make $\text{RMSD}(D_f, M) = 0$ by setting $f = 1/|D|$ and aligning any single point exactly to another point. So RMSD by itself is no longer a viable measure. For this reason, we propose the following distance measure.

Definition 2.2. [FRMSD] *The fractional root mean squared distance (or FRMSD) is defined as follows:*

$$\text{FRMSD}(D, M, f, \mu) = \frac{1}{f^\lambda} \sqrt{\frac{1}{|D_f|} \sum_{p \in D_f} \|p - \mu(p)\|^2}$$

We will empirically and mathematically justify a value of λ in Section 6.4 and Section 5. For brevity we sometimes drop μ and just write $\text{RMSD}(D, M)$ or $\text{FRMSD}(D, M, f)$.

Problem 2.2. [minimize FRMSD] *Compute the transformation $T \in \mathcal{T}$ and fraction $f \in [0, 1]$ to minimize:*

$$\min_{\substack{T \in \mathcal{T} \\ f \in [0, 1]}} \frac{1}{f^\lambda} \sqrt{\frac{1}{|D_f|} \sum_{p \in D_f} \|T(p) - \mu(p)\|^2}$$

Intuitively, the $\frac{1}{f^\lambda}$ term serves to balance the RMSD term.

3 Fractional ICP

Fractional ICP or FICP, detailed in Algorithm 3, finds a local minimum for FRMSD.

Algorithm 3.1 FICP(D, M)

-
- 1: Compute $\mu_0 = \arg \min_{\mu_0: D \rightarrow M} \text{RMSD}(D, M, \mu_0)$.
 - 2: Compute $f_0 \in [0, 1]$ minimizing $\text{FRMSD}(D, M, f_0, \mu_0)$.
 - 3: $i \leftarrow 0$.
 - 4: **repeat**
 - 5: Compute D_{f_i} minimizing $\text{RMSD}(D_{f_i}, M, \mu_i)$ such that $D_{f_i} \subseteq D$ and $|D_{f_i}| = \lfloor f_i |D| \rfloor$.
 - 6: Compute $T \in \mathcal{T}$ minimizing $\text{RMSD}(D_{f_i}, M, \mu_i)$. $D \leftarrow T(D)$.
 - 7: $i \leftarrow i + 1$.
 - 8: Compute $\mu_i : D \rightarrow M$ minimizing $\text{RMSD}(D, M, \mu_i)$.
 - 9: Compute $f_i \in [0, 1]$ minimizing $\text{FRMSD}(D, M, f_i, \mu_i)$.
 - 10: **until** ($u_i = u_{i-1}$ and $f_i = f_{i-1}$)
-

In practice, the comparison on line 10 of Algorithm 3 can be replaced by checking whether the $\text{FRMSD}(D, M, f)$ value decreases by less than some threshold.

3.1 Implementation

An implementation of ICP has two basic operations: computing the matching (which can be done efficiently with a kd-tree, a d^2 -tree [11], or with point-to-surface alignment [3, 16]) and computing the transformation (which can also be solved efficiently with a variety of approaches [10, 9, 19]). Extending such an implementation to FICP requires two more simple operations.

Computing the subset D_f . The set $D_f = \{p \in D \mid |D_f| = \lfloor f|D| \rfloor, \text{RMSD}(D_f, M) \text{ is minimized}\}$ is defined by the $\lfloor f|D| \rfloor$ points with the smallest residual distances $r = \|p - \mu(p)\|$. This observation implies the following algorithm. Compute and sort all residual distances and let D_f be the $f|D|$ points with the smallest residual distances. The runtime is bounded by the sorting which takes $O(|D| \log |D|)$ time.

Computing the fraction. There are only $|D|$ fractions which we need to consider. Consider the sorted order of the point set D by each point's residual distance $r = \|p - \mu(p)\|$. Each prefix of this ordering represents a distinct fraction. If we maintain the value $\sum_{p \in D_f} \|p - \mu(p)\|^2$ for each D_f we can compute $\text{FRMSD}(D, M, f)$ in constant time for a given fraction f . We can also update D_f to a point set of size $|D_f| + 1$ in constant time by adding the next point in the sorted order to D_f . If the i th prefix yields the smallest value of FRMSD , then f is set to $i/|D|$. So this computation takes $O(|D|)$ time.

3.2 Convergence of Algorithm

FICP converges to a local minimum of $\text{FRMSD}(D, M, f)$ in a sense that if all but one of transformations, matchings, or fractions is fixed, then the value of the remaining variable cannot be changed to decrease the value of $\text{FRMSD}(D, M, f)$.

Theorem 3.1. *For any two points sets $D, M \in \mathbb{R}^d$, Algorithm 3 converges to a local minimum of $\text{FRMSD}(T(D), M, f, \mu)$ over $(f, T, \mu) \in [0, 1] \times \mathcal{T} \times \{D \rightarrow M\}$.*

Proof Sketch: Algorithm 3 only changes the value of (f, T, μ) when computing the optimal transformation T (line 7), computing the optimal matching μ (line 8), or computing the optimal fraction f (line 9). None of these steps can increase the value of $\text{FRMSD}(D, M, f_i, \mu_i)$, because staying at the current value would retain the value of $\text{FRMSD}(D, M, f, \mu)$, but each can potentially decrease it. (Full proof appears in longer version [14].) \square

4 Data Generation Model

In order to formalize the expected mathematical properties of the FRMSD measure and the FICP algorithm, we now state some fairly general assumptions about the input data. All data on which FICP is used need not have these exact properties, but we hope that these properties are general enough that whatever differences exist in the alternative data will not significantly affect the following analysis and the resulting conclusions.

We assume that data points are generated from model points by the following abstract procedure:

1. Generate a set M_I of model points that will have corresponding data points
2. For every model point $m \in M_I$, let $p = T^{-1}(m + n)$ be the corresponding data point, where T is a transformation in the set \mathcal{T} and n is isotropic Gaussian noise with standard deviation σ . The set of data points p corresponding to M_I is denoted as D_I .
3. Generate a random set D_O of data outliers.
4. Generate a random set M_O of model outliers out of a spatial Poisson process.

We let $D = D_I \cup D_O$ and $M = M_I \cup M_O$. Let p_I be the fraction of data inliers relative to all data points. The Poisson process for model outliers is a minimally informative prior. We let the density of this process be ω points per unit volume.

The probability density of the squared magnitude $z = \|n\|^2$ of the correspondence noise is a chi square density in d dimensions:

$$g_{\chi^2(d)}(z) = \frac{z^{d/2-1}}{2^{d/2}\sigma^d\Gamma(d/2)} e^{-\frac{z}{2\sigma^2}}$$

where

$$\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$$

is the gamma function. The expected number of model outliers in a region of space with volume V is equal to ωV .

Suppose now that the correct geometric transformation $T \in \mathcal{T}$ is applied to data point p to obtain the transformed data point

$$q = T(p) = m + n.$$

If q and m correspond, their distance statistics are chi square. If q and m do not correspond, the situation is more complex: Either point (or both) could be an outlier, or they could be non-corresponding inliers. We do not know the distance statistics for model inliers. In the remainder of this section, we assume that the probability that a data inlier is nearest to a non-corresponding model inlier is negligible. Under this assumption, the probability density of the distance r from q to the nearest outlier, given that model outliers are from a spatial Poisson process with density ω points per unit volume, can be shown to be

$$w(r) = \omega S(d) r^{d-1} e^{-\omega S(d) r^d/d} \quad \text{for } r \geq 0$$

where

$$S(d) = \frac{2\pi^{d/2}}{\Gamma(d/2)}$$

is the surface of the unit sphere in d dimensions.

So far we have not specified the units of measure. Since σ is a distance and ω is a distance raised to power $-d$ (density per unit volume), the parameter $\sigma\omega^{1/d}$ is dimensionless. As long as σ and ω are properly scaled to each other, the analysis that follows is independent of σ .

5 The Value of λ

In this Section we justify a particular choice for the value of λ used in the definition of the fractional root mean squared distance (FRMSD).

As shown in Section 3.1, the FICP algorithm selects a fraction f of data-model matches in increasing order of their residual distances $r = \|p - \mu(p)\|$. Because of this, choosing a fraction f is equivalent to choosing a maximum allowed value r^* for the residual distance r . Since we would like the FICP algorithm to favor inliers over outliers, it makes sense to require r^* to be defined in such a way that data points that are r^* away from a model point are equally likely to be inliers as they are to be outliers. Let us call such a value of r^* the *critical distance*. We then ask the following question: *Is there a value of λ in the definition of the FRMSD for which the value of f chosen by the FICP algorithm corresponds to the critical distance?*

To answer this question, we first express r^* as a function of the model parameters (Section 5.1), and determine the function that relates an arbitrary distance r to the corresponding fraction f (Section 5.2). We then write an estimate of the FRMSD under an ergodicity assumption (Section 5.3). This estimate is itself a function of f , and therefore of r . The FICP algorithm maximizes the FRMSD with respect to f , that is, finds a zero for the derivative of the FRMSD with respect to f . Setting the value of f where this zero is achieved to $f(r^*)$ yields an equation for λ , whose solution set justifies our choice for this parameter (Section 5.4).

Our analysis holds for outlier densities ω that are below a certain value ω_{\max} , which is inversely proportional to the standard deviation σ of the noise that affects the data points. If outliers exceed this density, then matching data and model points based on minimum distance is too unreliable to yield good results.

5.1 The Critical Distance

The volume of a sphere of radius r in d dimensions is

$$V_s(r) = \frac{S(d)}{d} r^d$$

where $S(d)$ was defined in Section 4. The volume of the shell between radii r and $r + \delta r$ is

$$\delta V_s = \frac{S(d)}{d} [(r + \delta r)^d - r^d] \approx S(d) r^{d-1} \delta r.$$

This approximation is asymptotically exact as $\delta r \rightarrow 0$.

The probability mass in the same shell for an isotropic Gaussian distribution with zero mean and standard deviation σ is

$$\delta G_s = 2r g_{\chi^2(d)}(r^2) \delta r = \frac{S(d)}{(2\pi)^{d/2} \sigma} \left(\frac{r}{\sigma}\right)^{d-1} e^{-\frac{1}{2}\left(\frac{r}{\sigma}\right)^2} \delta r$$

as $\delta r \rightarrow 0$ (the term $2r$ derives from the Jacobian of the transformation $z = r^2$, since the χ^2 density is defined for the square of a distance).

Assume that the center of the shell above is at the transformed data point q defined in Section 4. As explained in Section 4, if q and m correspond, their distance statistics are chi squared, and the likelihood of a particular radius r is $\delta G_s / \delta r$. Otherwise, the distance statistics are approximately described by a spatial Poisson process with density ω . Then, the critical distance is determined by the equation

$$\omega \delta V_s = \delta G_s$$

that is,

$$e^{-\frac{1}{2}\left(\frac{r}{\sigma}\right)^2} = \omega \sigma^d (2\pi)^{d/2}. \quad (5.1)$$

The left-hand side of equation (5.1) is strictly positive and monotonically decreasing in r and the right-hand side is constant, so the equation admits a solution if and only if

$$0 < \omega \leq \omega_{\max} = \frac{1}{(\sqrt{2\pi} \sigma)^d}.$$

If the outliers exceed this maximum density ω_{\max} , the critical distance shrinks to zero: any model point around any given data point q is more likely to be an outlier than it is to be the model point corresponding to q . Of course, when there are no model outliers ($\omega = 0$) the concept of critical distance loses its significance.

Equation (5.1) can be solved for r to yield the desired value of r^* as a function of the model parameters:

$$\frac{r^*}{\sigma} = \sqrt{-2 \log_e((\sqrt{2\pi} \sigma)^d \omega)} = \sqrt{2 \log_e \frac{\omega_{\max}}{\omega}}.$$

5.2 Relationship between f and r

With probability p_I , the data point q has a corresponding model point (inlier). In this event, if r_I is the distance from this model point and r_O is the distance from the nearest model outlier point, the complement of the cumulative probability function of the distance r to the nearest model point (either inlier or outlier) is

$$\begin{aligned} 1 - F(r) &= 1 - \mathcal{P}[\min(r_I, r_O) < r] = \mathcal{P}[\min(r_I, r_O) \geq r] \\ &= \mathcal{P}[r_I \geq r \cap r_O \geq r] = \mathcal{P}[r_I \geq r] \mathcal{P}[r_O \geq r] \\ &= (1 - \mathcal{P}[r_I \leq r]) (1 - \mathcal{P}[r_O \leq r]) = (1 - F_I(r)) (1 - F_O(r)) \end{aligned}$$

where $F_I(r)$ and $F_O(r)$ are respectively the probability that the matching model point and the nearest model outlier are at most r units away from q . From Section 4, these probabilities are as follows:

$$F_I(r) = \int_0^{r^2} g_{\chi^2(d)}(\zeta) d\zeta \quad \text{and} \quad F_O(r) = \int_0^r w(\rho) d\rho.$$

Then, if q has a corresponding model point, the density of its distance from the nearest model point is

$$\begin{aligned} \phi_c(r) &= \frac{dF(r)}{dr} = -\frac{d}{dr}(1 - F(r)) \\ &= 2r g_{\chi^2(d)}(r^2) (1 - F_O(r)) + (1 - F_I(r)) w(r). \end{aligned}$$

With probability $p_O = 1 - p_I$, the data point q is instead an outlier. Then, it has no corresponding model point, so the probability that the nearest model point is at most r units away is simply $F_O(r)$. In summary, the probability density of the distance between a data point q and its nearest model point $\mu(q)$ is

$$\phi(r) = p_I \phi_c(r) + p_O w(r)$$

and the average fraction of model points within r units from a data point is

$$f(r) = \int_0^r \phi(\rho) d\rho = p_I \int_0^r \phi_c(\rho) d\rho + p_O F_O(r).$$

The derivative of f with respect to r is $\phi(r)$.

5.3 Ergodic Estimate of the FRMSD

An estimate of the fractional root mean squared distance (FRMSD) can be obtained by assuming ergodically that the sample moment included in the definition of FRMSD is close to the corresponding statistical moment:

$$\frac{1}{f|D|} \sum_{p \in D_f} \|p - \mu(p)\|^2 \approx \mathbf{E}_{p \in D_f} [\|p - \mu(p)\|^2].$$

We can then write

$$\begin{aligned} \text{FRMSD}^2(D, M, f) &= \frac{1}{f^{2\lambda}} \frac{1}{f|D|} \sum_{p \in D_f} \|p - \mu(p)\|^2 \\ &\approx \frac{1}{f^{2\lambda}} \mathbf{E}_{p \in D_f} [\|p - \mu(p)\|^2] = \frac{1}{f^{2\lambda}} \int_0^r \rho^2 \phi(\rho) d\rho. \end{aligned}$$

5.4 Stationary Point of the FRMSD Estimate

At the minimum of $\text{FRMSD}(D, M, f)$, the derivative of $\text{FRMSD}^2(D, M, f)$ with respect to f is zero. Differentiation of the expression at the end of Section 5.3 yields

$$\frac{d}{df} \text{FRMSD}^2(D, M, f) = \frac{-2\lambda}{f^{2\lambda+1}} \int_0^r \rho^2 \phi(\rho) d\rho + \frac{r^2}{f^{2\lambda}} \phi(r) \frac{dr}{df}$$

Since

$$\left(\frac{dr}{df}\right)^{-1} = \frac{df}{dr} = \phi(r),$$

the last addend simplifies to $r^2/f^{2\lambda}$, and

$$f^{2\lambda} \frac{d}{df} \text{FRMSD}^2(D, M, f) = -\frac{2\lambda}{f} \int_0^r \rho^2 \phi(\rho) d\rho + r^2.$$

Zeroing this derivative and setting $r = r^*$ and $f = f(r^*)$ yields the following equation for λ :

$$\lambda = \frac{1}{2} \frac{(r^*)^2 \int_0^{r^*} \phi(\rho) d\rho}{\int_0^{r^*} \rho^2 \phi(\rho) d\rho}.$$

Figure 1 plots the values of λ in two and three dimensions as a function of the relative model outlier density ω/ω_{\max} and for various values of the data inlier fraction p_I .




Figure 1. Theoretical value of λ in the definition of the FRMSD in two (upper bundle) and three (lower bundle) dimensions as a function of the relative model outlier density ω/ω_{\max} . Curves in each bundle correspond to $p_I = \{0.5, 0.6, 0.7, 0.8, 0.9\}$ from the bottom up. Dependency on p_I is weak.

Since the noise standard deviation σ acts merely as an overall scale factor, these plots do not depend on σ . It is apparent from the figure that λ depends weakly on the fraction p_I of data inliers. The knees of the plots are at about $\lambda = 1.3$ and $\lambda = 0.95$ for $d = 2$ and $d = 3$ dimensions, respectively, corresponding to $\omega/\omega_{\max} = 0.2$. These knee values are selected as general-purpose values for the definition of FRMSD in two and three dimensions.

6 Experiments

The main advantage of FICP over other variants of ICP is that it automatically determines the outlier set via a fraction

f and reaches a optimum in terms of the correspondence, the transformation, and the fraction of outliers. In doing so, it takes less time than algorithms which have no guarantees, despite searching a larger parameter space.

We deal empirically with the issue of the parameter λ used in the definition of FRMSD. We observe that FRMSD is robust to the choice of λ within a broad range. However the radius of convergence and efficiency of FICP is improved when λ is set to a slightly higher values than those determined optimal for identifying outliers in Section 5. Intuitively, a smaller value of λ is more likely to classify correct correspondences as outliers when the alignment is not close, and thus get stuck in local minimum. For higher values of λ these types of local minimum seem less prevalent. So for all performance studies we set $\lambda = 3$, unless otherwise specified. For this value FICP has an expanded radius of convergence and tends to find very similar alignments as when λ is set according to the analysis in Section 5. After converging, we recommend setting $\lambda = 1.3$ for $d = 2$ or $\lambda = .95$ for $d = 3$ to identify outliers more aggressively. This final phase should take very few additional iterations of the algorithm, since, as we demonstrate, moderately modifying the value of λ has small effects on the FRMSD and f values returned.

6.1 Data Sets

We use the SQUID fish contour database [18] from the University of Surrey, UK. This database has 1100 2D contours of fish and each contour has 500 to 3000 points. The size of this data set allows us to average results over a very large set of experiments.

We also perform some experiments on a limited number of 3D models. In particular we use the *bunny* (35,947 points) and the *happy Buddha* (144,647 points) data set from the Stanford 3D Scanning Repository.

We synthetically introduce outliers into the data sets in 3 ways. We always begin by creating two copies M and D , to represent the model and the input data, of the particular data set. A parameter p_I fraction of the final set D are left undisturbed as data inliers.

- **Occlusion:** We randomly choose a ball B and remove all of the points from M within B . This represents cases where the model set is only partially observed because of occlusions, where two overlapping views of the same object do not exactly align, or where the input data D has grown since the model was formed.
- **Deform:** We randomly choose a ball B and shift randomly the points $D \cap B$. This represents the case where D is deformed slightly between time steps.
- **New data:** We add a set of points to D . These points are placed uniformly at random within a bounding box

of D . This represents outliers caused by some sort of data retrieval noise or from spurious or new data.

Finally, we independently add Gaussian noise to each point $p \in D$.

We perform many tests on synthetic data because we know that a good match exists and it is thus easy to quantify the performance on our algorithm.

Additionally, we perform tests on real scanned data. We align pairs of scanned images of the *dragon* model (*dragonStandRight*) from the Stanford 3D Scanning Repository from views 24° or 48° apart. Because the different views observe different portions of the model, there are many points which have no good alignment in both the model and data set. These are outliers.

6.2 Performance

For each data set and type of outliers described above, we perform the following set of tests. Results are averaged over all SQUID data sets or 10 random rotations for 3D models. We first rotate D by θ degrees where θ is from the set $\{5^\circ, 10^\circ, 25^\circ, 50^\circ\}$. The axis of rotation is chosen randomly for the 3D data. We then run ICP, TrICP searching for f with the golden section search [4], and FICP, minimizing over all rigid motions. We report the total number of iterations of each, the run time, and the final values of RMSD, FRMSD, and f . We vary the input so that p_I is either $\{.75, .88, .95\}$. We expect that optimally f should be near p_I since in our data ω/ω_{\max} is small. All experiments were performed on a 3 GHz Pentium IV processor with 1 Gb SD-RAM.

Tables 1, 2, 3, and 4 show a sample of these results. TrICP and FICP return similar values of RMSD and FRMSD on average while also determining reasonable values for f . However, FICP is about $6 \times$ to $11 \times$ faster than TrICP using the golden section search.

Alg.	p_I	time (s)	# iter.	RMSD	FRMSD	f
ICP	.75	0.335	24.5	9.454	9.454	1.000
TrICP	.75	1.356	117.9	0.217	0.541	0.744
FICP	.75	0.178	13.6	0.178	0.424	0.749

Table 1. SQUID: Occlusion outliers, rotated 5° .

Alg.	p_I	time (s)	# iter.	RMSD	FRMSD	f
ICP	.75	0.461	26.7	5.820	5.820	1.000
TrICP	.75	1.578	92.9	0.176	0.399	0.768
FICP	.75	0.264	13.7	0.175	0.388	0.766

Table 2. SQUID: New Data outliers, rotated 5° .

Observe in Figure 2 how in the alignment of the bunny data set, the non-deformed points (red points on back side,

Alg.	p_I	time (s)	# iter.	RMSD	FRMSD	f
ICP	.88	29.6	48.0	0.4530	0.4530	1.000
TrICP	.88	147.1	224.3	0.0052	0.0077	0.880
FICP	.88	13.7	15.9	0.0052	0.0077	0.880

Table 3. bunny: Deform outliers, rotated 5°

Alg.	p_I	time (s)	# iter.	RMSD	FRMSD	f
ICP	.88	109.2	28.2	0.2975	0.2975	1.000
TrICP	.88	485.4	120.5	0.0012	0.0018	0.880
FICP	.88	81.4	15.7	0.0012	0.0017	0.880

Table 4. Buddha: Deform outliers, rotated 5°

blue points are not visible because they lie exactly behind the red points) are aligned almost exactly by the FICP algorithm while the deformed points (shifted from visible blue points in front) are ignored. Such an alignment allows one to easily identify the portion of the data which has been deformed, and by how much it has been deformed. Without a proper registration to the model, the unaligned points have no point of comparison to gauge their deformation. The alignment is skewed when ICP is used and it is not helpful in determining which points are deformed.




Figure 2. bunny: M in blue (top left) and D in red (top right) with Deform noise with $p_I = .75$. Registered using FICP (bottom left) and ICP (bottom right).

6.3 Funnel of Convergence

We measure the radius of convergence of each algorithm by calculating the percentage of cases from the SQUID data

set that converge to an FRMSD value within .01 and f value within .01 of the alignment between the same sets with no initial rotation. Table 5 shows the results when New Data outliers with $p_I = .88$ are added to the data set D . The results for the other types of noise are similar. For 3D data sets we chose σ proportionally smaller, so these convergence rates are all slightly larger. Note that FICP with $\lambda = 3$ performs much better than when $\lambda = 1.3$.

Alg.	λ	5°	10°	25°	50°
ICP	-	0.999	0.997	0.994	0.962
TrICP	3	0.875	0.870	0.853	0.816
FICP	3	0.952	0.945	0.909	0.875
FICP	1.3	0.857	0.473	0.141	0.060

Table 5. Percentage of SQUID data sets converging per initial rotation.

ICP has a larger radius of convergence than FICP, because it searches a much smaller parameter space. FICP has a larger radius of convergence than TrICP even though they search the same parameter space.

6.4 Validating λ

We empirically justify that FRMSD is not sensitive to the choice of λ . We run FICP with λ set to $\{1, 1.3, 2, 3, 4, 5\}$. We plot the averaged results on the SQUID data set when Occlusion noise is added with $p_I = .75$ and D is initially rotated 0° and 5° in Table 6 and Table 7, respectively. Altering λ does not dramatically affect the converged solution, but can affect the radius of convergence. The output is similar for different types of noise. On 3D data, FICP performs slightly better than 2D data for smaller λ .

Alg.	λ	time (s)	# iter.	RMSD	FRMSD	f
FICP	1	0.142	10.38	0.158	0.225	0.701
FICP	1.3	0.069	3.81	0.170	0.248	0.749
FICP	2	0.059	3.06	0.170	0.303	0.750
FICP	3	0.061	3.17	0.170	0.404	0.750
FICP	4	0.062	3.21	0.171	0.538	0.751
FICP	5	0.063	3.30	0.172	0.717	0.751

Table 6. FICP, varying λ , with D rotated 0° .

6.5 Aligning Scanned Model Data

We aligned the *dragon* scans 24° and 48° apart with ICP, TrICP, and FICP. See longer version [14] for tables. For most alignments both FICP and TrICP realize an alignment with a much lower FRMSD value than ICP, with FICP, occasionally, noticeably outperforming TrICP. FICP is usually

Alg.	λ	time (s)	# iter.	RMSD	FRMSD	f
FICP	1	0.733	37.23	0.298	1.503	0.274
FICP	1.3	0.488	36.44	0.219	0.563	0.660
FICP	2	0.244	17.00	0.176	0.329	0.740
FICP	3	0.198	13.59	0.178	0.424	0.749
FICP	4	0.194	13.28	0.184	0.570	0.751
FICP	5	0.200	13.66	0.299	0.875	0.756

Table 7. FICP, varying λ , with D rotated 5° .

about as fast as ICP, and is consistently about $5\times$ to $10\times$ faster than TrICP.

Figure 3 shows the alignment of the scan at 0° aligned with the scan at 48° using ICP and FICP. Notice how when the scans are aligned with ICP, the points in the dragon's tail are slightly misaligned, whereas with FICP, the alignment is much better.




Figure 3. Dragon scans at 0° and 48° with M in blue and D in red, registered using FICP (top left) and ICP (top right). Zoomed images of the alignment around dragon's tail with FICP (bot. left) and ICP (bot. right) demonstrate the skew in the alignment due to ICP.

References

[1] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. on Pat. Anal. and Mach. Intel.*, 14(2), February 1992.

[2] C.-S. Chen, Y.-P. Hung, and J.-B. Cheng. RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. on Pat. Anal. and Mach. Intel.*, 21(11), November 1999.

[3] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Im. and Vis. Comp.*, 10:145–155, 1992.

[4] D. Chetverikov, D. Stepanov, and P. Krsek. Robust euclidean alignment of 3d point sets: the trimmed iterative closest point algorithm. *Im. and Vis. Comp.*, 23(3):299–309, 2005.

[5] G. Dalley and P. Flynn. Pair-Wise Range Image Registration: A Study of Outlier Classification. *Comp. Vis. and Im. Und.*, 87:104–115, 2002.

[6] N. Gelfand, L. Ikemoto, S. Rusinkiewicz, and M. Levoy. Geometrically stable sampling for the icp algorithm. In *3DIM*, 2003.

[7] N. Gelfand, M. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global alignment. *EG Symp. Geom. Proc.*, 2005.

[8] W. E. Grimson, R. Kikinis, F. A. Jolesz, and P. M. Black. Image-guided surgery. *Sci. Am.*, 280:62–69, 1999.

[9] R. J. Hanson and M. J. Norris. Analysis of measurements based on the singular value decomposition. *SIAM J. Sci. and Stat. Comp.*, 27(3):363–373, 1981.

[10] B. K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. of Am. A*, 4, April 1987.

[11] S. Leopoldseder, H. Pottmann, and H.-K. Zhao. The d^2 -tree: A hierarchical representation of the squared distance function. Technical Report 101, Inst. of Geom., Vienna University of Technology, 2003.

[12] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The Digital Michelangelo Project: 3D Scanning of Large Statues. *AMC SIGGRAPH*, 2000.

[13] X. Li and I. Guskov. Multi-scale features for approximate alignment of point-based surfaces. In *EG Symp. Geom. Proc.*, 2005.

[14] J. M. Phillips, R. Liu, and C. Tomasi. Outlier robust ICP for minimizing fractional RMSD. *ArXiv*, (cs.GR/0606098), June 2006.

[15] H. Pottmann, Q.-X. Huang, Y.-L. Yang, and S.-M. Hu. Geometry and convergence analysis of algorithms for registration of 3D shapes. Technical Report 117, Geometry Preprint Series, TU Wien, June 2004.

[16] S. Rusinkiewicz and M. Levoy. Efficient variants of the ICP algorithm. *3DIM*, 2001.

[17] G. Turk and M. Levoy. Zippered Polygonal Meshes from Range Images. *ACM SIGGRAPH*, 1994.

[18] U. University of Surry, Guilford. The squid database: Shape queries using image databases.

[19] T. D. Wu, S. C. Schmidler, T. Hastie, and D. L. Brutlag. Modeling and superposition of multiple protein structures using affine transformations: Analysis of the globins. *Pac. Sym. on Bio.*, pages 507–518, 1998.

[20] Z. Zhang. Iterative point matching for registration of free-form curves. *Int. J. Comp. Vis.*, 7(3):119–152, 1994.