Iterative Most Likely Oriented Point Registration

Seth Billings and Russell Taylor

Johns Hopkins University, Department of Computer Science, Baltimore, MD, USA
{sbillin3,rht}@jhu.edu

Abstract. A new algorithm for model based registration is presented that optimizes both position and surface normal information of the shapes being registered. This algorithm extends the popular Iterative Closest Point (ICP) algorithm by incorporating the surface orientation at each point into both the correspondence and registration phases of the algorithm. For the correspondence phase an efficient search strategy is derived which computes the most probable correspondences considering both position and orientation differences in the match. For the registration phase an efficient, closed-form solution provides the maximum likelihood rigid body alignment between the oriented point matches. Experiments by simulation using human femur data demonstrate that the proposed Iterative Most Likely Oriented Point (IMLOP) algorithm has a strong accuracy advantage over ICP and has increased ability to robustly identify a successful registration result.

Keywords: point cloud registration, Fisher distribution, PD tree.

1 Introduction

The need to register multiple representations of anatomy is a problem frequently encountered in the medical imaging and computer assisted intervention domains. As an example, computer assisted total hip replacement may involve sampling points from a femur bone during surgery and registering those points with a surface model derived from preoperative CT images. For such clinical applications, the Iterative Closest Point (ICP) algorithm has been extensively applied with many variants [1] since its introduction by Besl and McKay [2]. Surface normals could also be acquired in this context from range imaging techniques, surface probing, etc., enabling use of points and orientations as proposed in this paper.

1.1 Iterative Closest Point (ICP) Algorithm

Consider a source shape represented by a set of n points $X = \{\boldsymbol{x_i}\}$ and a target shape represented by Ψ (typ. another point cloud or a mesh). The ICP algorithm seeks to compute the rigid body transformation T that minimizes the sum of square distances between the two shapes

$$T = \underset{T}{\operatorname{argmin}} \sum_{i=1}^{n} \|\boldsymbol{y_i} - T(\boldsymbol{x_i})\|_2^2 \quad (1)$$

P. Golland et al. (Eds.): MICCAI 2014, Part I, LNCS 8673, pp. 178–185, 2014.

where $\boldsymbol{y_i}$ is the point on the target shape closest to $T(\boldsymbol{x_i})$ as defined by the closest point correspondence operator

$$\boldsymbol{y_i} = \mathrm{C_{CP}}(T(\boldsymbol{x_i}), \Psi) = \underset{\boldsymbol{y} \in \Psi}{\operatorname{argmin}} \|\boldsymbol{y} - T(\boldsymbol{x_i})\|_2 \ . \tag{2}$$

The ICP algorithm may be summarized as an iteration of two key steps:

1. Compute correspondences between the source and target shapes.
2. Compute the transformation T that minimizes the sum of square distances between the correspondences.

The first step has an efficient implementation using a KD tree search, while the second step has a closed form solution via Arun's method [3].

In this paper we present a novel variant of ICP that incorporates surface orientations alongside position data. In [4] surface normals were considered in the correspondence phase of an ICP method by limiting match orientation differences to within 45 degrees. Lara et al. [5] extended this approach by adaptively computing allowable bounds on the match orientation errors. In [6] contour normals were used in an expectation maximization (EM) framework for 3D-2D registration involving single X-ray images. Other ICP variants are found in [1].

Rather than using orientations to narrow the range of permitted matches, as in [4,5], our method combines orientation and position information in a cohesive probabilistic framework within both the correspondence and registration phases of the algorithm in order to directly compute the most probable matches and optimize match alignment. Thus, we name our method the Iterative Most Likely Oriented Point (IMLOP) algorithm. Our probabilistic model is similar to that of [6], though [6] is not a 3D-3D method. For the correspondence phase, we devise a novel tree search strategy that efficiently computes the most likely oriented point correspondences. We also incorporate a closed-form solution to the match alignment subproblem of the registration phase that optimizes alignment of both match positions and orientations, which was previously addressed in [7]. Our methods are summarized in Algorithms 1-3 to follow.

2 Methods

In this section, we present the proposed IMLOP algorithm, beginning with an overview of the method and then describing in detail the computations of the correspondence and registration sub-phases of the algorithm. For the purposes of this section, consider a source shape represented by a set of n oriented sample points $X = \{(\boldsymbol{x}_{\mathrm{p}\boldsymbol{i}}, \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}})\}$ where $\boldsymbol{x}_{\mathrm{p}\boldsymbol{i}}$ is a position vector and $\boldsymbol{x}_{\mathrm{n}\boldsymbol{i}}$ is an orientation unit vector associated with oriented sample point $\boldsymbol{x_i}$. Also consider a target shape represented by Ψ, such as a surface mesh or a second oriented point cloud.

2.1 Iterative Most Likely Oriented Point (IMLOP) Algorithm

The IMLOP algorithm incorporates a probabilistic framework formulated using Fisher and Gaussian distributions to model the measurement errors of orientation and position, respectively. Since the Fisher distribution is the analogue of

Gaussian on the unit sphere, pairing them to model oriented point measurement error is both natural and analytically convenient. Assuming unbiased, iid error and independence of orientation and position, the PDF describing the probability that a measured oriented point $\boldsymbol{x} = (\boldsymbol{x}_{\mathrm{p}}, \boldsymbol{x}_{\mathrm{n}})$ corresponds to $\boldsymbol{y} = (\boldsymbol{y}_{\mathrm{p}}, \boldsymbol{y}_{\mathrm{n}})$ has the form

$$f_{\mathrm{match}}(\boldsymbol{x}, \boldsymbol{y}) = \frac{k}{(2\pi\sigma^2)^{3/2} \cdot 2\pi(e^k - e^{-k})} \cdot e^{k\boldsymbol{y}_{\mathrm{n}}{}^T \boldsymbol{x}_{\mathrm{n}} - \frac{1}{2\sigma^2}\|\boldsymbol{y}_{\mathrm{p}} - \boldsymbol{x}_{\mathrm{p}}\|_2^2} \tag{3}$$

where k is a concentration parameter for orientation error and σ^2 is the variance of positional error. The correspondence phase of IMLOP selects the most probable match for each $T(\boldsymbol{x_i})$ via the most likely point correspondence operator

$$\boldsymbol{y_i} = \mathrm{C}_{\mathrm{MLP}}(T(\boldsymbol{x_i}), \Psi) = \underset{\boldsymbol{y} \in \Psi}{\operatorname{argmax}}\, f_{\mathrm{match}}(T(\boldsymbol{x_i}), \boldsymbol{y}) \tag{4}$$

and the registration phase solves a new transformation T to maximize the correspondence likelihood over all point pairs, which reduces to solving

$$T = \underset{T}{\operatorname{argmin}} \left(\frac{1}{2\sigma^2} \sum_{i=1}^{n} \|\boldsymbol{y}_{\mathrm{p}\boldsymbol{i}} - T(\boldsymbol{x}_{\mathrm{p}\boldsymbol{i}})\|_2^2 - k \sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}{}^T R \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} \right) \tag{5}$$

where R is the rotation component of T. The IMLOP algorithm is summarized in Algorithm 1.

Algorithm 1. Iterative Most Likely Oriented Point (IMLOP)

input : Source shape X and target shape Ψ
Initial noise parameters k_0 and σ_0^2
Initial transformation T_0
output: Final transformation T that aligns X and Ψ

```
Initialize: T ← T_0, k ← k_0, σ^2 ← σ_0^2
while not converged do
    Compute oriented point correspondences: y_i = C_MLP(T(x_i), Ψ), i = 1..n
    Register oriented point correspondences:
    T = argmin_T ( 1/(2σ^2) Σ_{i=1}^{n} ||y_pi − T(x_pi)||_2^2 − k Σ_{i=1}^{n} y_ni^T R x_ni )
    Update noise parameters: k = R̄(3−R̄^2)/(1−R̄^2), σ^2 = 1/n Σ_{i=1}^{n} ||y_pi − T(x_pi)||_2^2
end
```

The noise parameters k and σ^2 are estimated from the residual match errors at each iteration. The variance of position error (σ^2) is simply estimated as the mean square distance between matches. The concentration of orientation error (k) is estimated by an approximation to its maximum likelihood estimate [8,9]

$$k \approx \frac{\bar{R}(3 - \bar{R}^2)}{1 - \bar{R}^2}, \quad \bar{R} = \frac{1-w}{n} \sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}{}^T R \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} + \frac{w}{\alpha} \sum_{i=1}^{n} \boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}}{}^T R \boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}}, \tag{6}$$

$$\boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}} = \boldsymbol{x}_{\mathrm{p}\boldsymbol{i}} - \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{x}_{\mathrm{p}\boldsymbol{i}}, \quad \boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}} = \boldsymbol{y}_{\mathrm{p}\boldsymbol{i}} - \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{p}\boldsymbol{i}}, \quad \alpha = \sum_{i=1}^{n} (\|\boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}}\| \cdot \|R\boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}}\|) .$$

For a continuous surface, it is often possible to find a nearly perfect orientation match for any given orientation. Thus, estimating k based on the matched orientation differences alone may tend to progressively over-estimate k's value until positional differences would no longer be relevant. To prevent this, we consider the rotational misalignment represented in both match orientations and match positions when computing $\bar{R}$ in (6), which balances k at reasonable values. In our implementation, we equally weight each term using $w = 0.5$. In light of the above, one may also wish to restrict the effect of position errors to only decrease orientation confidence. In this case, the following alternative for $\bar{R}$ may be used

$$\bar{R} = \min\left(\frac{1}{n}\sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}{}^{T} R \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}}\,,\, \frac{1-w}{n}\sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}{}^{T} R \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} + \frac{w}{\alpha}\sum_{i=1}^{n} \boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}}{}^{T} R \boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}}\right) \quad . \tag{7}$$

2.2 Computing Most Likely Oriented Point Correspondences

This section describes an efficient method for implementing the most likely point correspondence operator of (4). Our implementation is based on a modified principal direction (PD) tree search, though the method is equally suited to modifying the more standard KD tree. The idea is to construct a tree search structure around the target shape using positional data in the standard manner. During tree construction, a minimal bounding box enclosing the shape is computed for each node. The average surface orientation ($\boldsymbol{N}_{\mathrm{avg}}$) and the maximum angular deviation from the average orientation ($\theta_{\max}$) is also computed within each node.

$$\boldsymbol{N}_{\mathrm{avg}} = \frac{\sum_{j\in Node} \boldsymbol{y}_{\mathrm{n}\boldsymbol{j}}}{\|\sum_{j\in Node} \boldsymbol{y}_{\mathrm{n}\boldsymbol{j}}\|}\,, \quad \theta_{\max} = \max_{j\in Node}\{\mathrm{acos}\,(\boldsymbol{N}_{\mathrm{avg}}{}^{T}\boldsymbol{y}_{\mathrm{n}\boldsymbol{j}})\} \tag{8}$$

Equivalent to maximizing the match probability of (3), the most likely point correspondence operator computes a match to minimize the match error equation

$$\mathrm{E}_{\mathrm{match}}(\boldsymbol{x},\boldsymbol{y}) = \frac{1}{2\sigma^2}\|\boldsymbol{y}_{\mathrm{p}} - \boldsymbol{x}_{\mathrm{p}}\|_2^2 + k(1 - \boldsymbol{y}_{\mathrm{n}}{}^{T}\boldsymbol{x}_{\mathrm{n}}) \tag{9}$$

which is always positive, unlike the form in (5).

When performing a correspondence search for a given oriented point $\boldsymbol{x}$, we begin with a current guess for the best correspondence (such as the match from the previous iteration), which has some match error E_{best}. At each node in the tree we require a fast test whether any point within the node may produce a lower match error than the current guess. Noting that lower match error is obtained for smaller orientation differences, we compute a lower bound on the difference in orientation between $\boldsymbol{x}_{\mathrm{n}}$ and any orientation $\{\boldsymbol{y}_{\mathrm{n}\boldsymbol{j}}\}$ within the node as

$$\theta_{\min} = \max\left(0\,,\, \mathrm{acos}\left(\boldsymbol{N}_{\mathrm{avg}}{}^{T}\boldsymbol{x}_{\mathrm{n}}\right) - \theta_{\max}\right) \quad . \tag{10}$$

It follows that for match errors lower than E_{best} there is an upper bound on the positional match distance given by

$$d_{\max} = \sqrt{2\sigma^2\left[E_{\mathrm{best}} - k(1 - \cos(\theta_{\min}))\right]} \quad . \tag{11}$$

Thus, if $\boldsymbol{x}_{\mathrm{p}}$ lies at a distance greater than $d_{\max}$ from a node bounding box, then the node must not contain a better match and may be skipped. This PD tree search strategy is summarized in Algorithm 2.

Algorithm 2. PD Tree Node Search for Oriented Point Correspondences

input : Oriented source point: $\boldsymbol{x}$
Current noise parameters: k and σ^2
Best match so far: $\boldsymbol{y}_{\mathrm{best}}$, E_{best}
This node object: $\mathcal{N}$
output: Updated best match: $\boldsymbol{y}_{\mathrm{best}}$, E_{best}

1 Compute $\theta_{\min}$ and $d_{\max}$ for this node
2 $\mathcal{B} \leftarrow$ bounding box of $\mathcal{N}$ expanded by $d_{\max}$ in all directions
3 **if** $\mathcal{B}$ *contains* $\boldsymbol{x}_p$ **then**
4   **if** $\mathcal{N}$ *is a leaf node* **then**
5   **foreach** $\boldsymbol{y_j} \in \mathcal{N}$ **do**
6   Compute match error: $\mathrm{E}_j \leftarrow \mathrm{E}_{\mathrm{match}}(\boldsymbol{x}, \boldsymbol{y_j})$
7   **if** $\mathrm{E}_j < E_{best}$ **then** Update best match: $\boldsymbol{y}_{\mathrm{best}} \leftarrow \boldsymbol{y_j}$, $E_{\mathrm{best}} \leftarrow \mathrm{E}_j$
8   **end**
9   **else**
10   Search left and right child nodes
11   **end**
12 **else** Skip node

Note that when the target shape is a mesh (rather than point cloud) the tree is constructed by representing each triangle in the mesh as a single point (such as the triangle center). After building the tree, the node bounding boxes are expanded as necessary to fully enclose all vertices of the triangles represented. Then in line 6 of Algorithm 2 the procedure for searching a leaf node is to first compute the point on the triangle represented by $\boldsymbol{y_j}$ that is closest to $\boldsymbol{x}_{\mathrm{p}}$ and then incorporate the triangle normal to compute the full match error.

2.3 Registering Oriented Point Correspondences

This section presents a closed-form solution to the subproblem of computing the rigid body transformation T, composed of a rotation R and translation t, that solves the minimization of (5). This minimization turns out to be a special case of a more general form and solution proposed in [7]. For completeness sake, we present the solution for this special case as Algorithm 3.

Recentering the point positions (line 1 of Algorithm 3) enables factoring the translation component of T out of (5). Rotation R is then found by minimizing

$$\begin{aligned} E &= \frac{1}{2\sigma^2} \sum_{i=1}^{n} \|\boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}} - R\boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}}\|_2^2 - k \sum_{i=1}^{n} {\boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}}^T R\boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} \qquad (12) \\ &= \frac{1}{2\sigma^2} \left[\sum_{i=1}^{n} \|\boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}}\| - 2\sum_{i=1}^{n} {\boldsymbol{y}'_{\mathrm{p}\boldsymbol{i}}}^T R\boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}} + \sum_{i=1}^{n} \|R\boldsymbol{x}'_{\mathrm{p}\boldsymbol{i}}\| \right] - k \sum_{i=1}^{n} {\boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}}^T R\boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} \end{aligned}$$

which is equivalent to maximizing

$$F = \frac{1}{\sigma^2} \sum_{i=1}^{n} \boldsymbol{y'}_{\mathrm{p}\boldsymbol{i}}{}^T R \boldsymbol{x'}_{\mathrm{p}\boldsymbol{i}} + k \sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}{}^T R \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} \tag{13}$$

which is solved by modifying the SVD method of Arun [3] as shown in line 2.

Algorithm 3. Register Oriented Point Correspondences

input : Corresponding oriented point sets: $X = \{\boldsymbol{x_i}\}$ and $Y = \{\boldsymbol{y_i}\}$
Noise parameters: k and σ^2

output: Transformation T that aligns X and Y

1 Compute new positions: $\boldsymbol{x'}_{\mathrm{p}\boldsymbol{i}} \leftarrow \boldsymbol{x}_{\mathrm{p}\boldsymbol{i}} - \bar{\boldsymbol{x}}_{\mathrm{p}}$, $\boldsymbol{y'}_{\mathrm{p}\boldsymbol{i}} \leftarrow \boldsymbol{y}_{\mathrm{p}\boldsymbol{i}} - \bar{\boldsymbol{y}}_{\mathrm{p}}$
where $\bar{\boldsymbol{x}}_{\mathrm{p}} = \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{x}_{\mathrm{p}\boldsymbol{i}}$, $\bar{\boldsymbol{y}}_{\mathrm{p}} = \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{y}_{\mathrm{p}\boldsymbol{i}}$

2 Compute covariance matrix: $H \leftarrow kH_1 + \frac{1}{\sigma^2} H_2$
where $H_1 = \sum_{i=1}^{n} \boldsymbol{x'}_{\mathrm{p}\boldsymbol{i}} \boldsymbol{y'}^{T}_{\mathrm{p}\boldsymbol{i}}$, $H_2 = \sum_{i=1}^{n} \boldsymbol{x}_{\mathrm{n}\boldsymbol{i}} \boldsymbol{y}_{\mathrm{n}\boldsymbol{i}}{}^T$

3 Compute rotation from SVD of H: $R \leftarrow VU^T$ where $H = USV^T$

4 **if** $\det(R) = -1$ **then** $R \leftarrow V'U^T$ where $V = [\boldsymbol{v_1}, \boldsymbol{v_2}, \boldsymbol{v_3}]$, $V' = [\boldsymbol{v_1}, \boldsymbol{v_2}, -\boldsymbol{v_3}]$

5 Compute translation: $t \leftarrow \bar{\boldsymbol{y}}_{\mathrm{p}} - R\bar{\boldsymbol{x}}_{\mathrm{p}}$

6 $T \leftarrow [R, t]$

3 Experiments

In this section, a simulation study (motivated by the introductory clinical scenario) is presented using a human femur segmented from CT imaging to evaluate the performance of IMLOP relative to ICP under known ground truth. A common termination condition is applied for both ICP and IMLOP, requiring the magnitude of change in the transformation T to be less than 0.001 mm and 0.001 degrees for two consecutive iterations. A triangular mesh of the femur surface is used as the target shape (Fig. 1). The source shape is constructed by randomly sampling oriented points from a subregion of the target shape (identified as a dark patch in Fig. 1). Registrations are performed after adding random noise to the oriented point positions and orientations and after applying a random misalignment selected uniformly from $[10, 20]$ degrees and $[10, 20]$ millimeters. Registrations are performed for different source shape sample sizes of 10, 20, 35, 50, 75 and 100 points and for different noise levels of 0, 0.5, 1 and 2 millimeters (degrees) standard deviation of Gaussian (wrapped Gaussian) noise applied to the sample positions (orientations). Three-hundred randomized trials are performed for each sample size / noise level pair. The accuracy of registration is evaluated using 100 validation points selected randomly from the mesh, with the average distance between registered and ground truth position forming the positional target registration error (TRE). An orientation TRE is similarly defined.

Figure 2 compares "success" TRE results and "failure" rates of ICP and IMLOP for various sample sizes at two noise levels. Registration failures are automatically detected using a threshold on the average of final match distances. For

IMLOP, the failure test is extended by an "or" condition to include a threshold on the final match orientation errors as well, enabling a more robust determination of registration success as may be appreciated by inspection of Fig. 1, which provides a detailed look at all registration trials for one sample size and noise level. In order to better appreciate the relationship between the final match errors and true registration error for each algorithm, trials along the x-axis in Fig. 1 are sorted by their positional TRE. For Fig. 2, we have set the match error threshold for registration failure at twice the standard deviation of applied noise. Plots of the remaining two noise levels and plots using different match error thresholds to flag success show similar results.

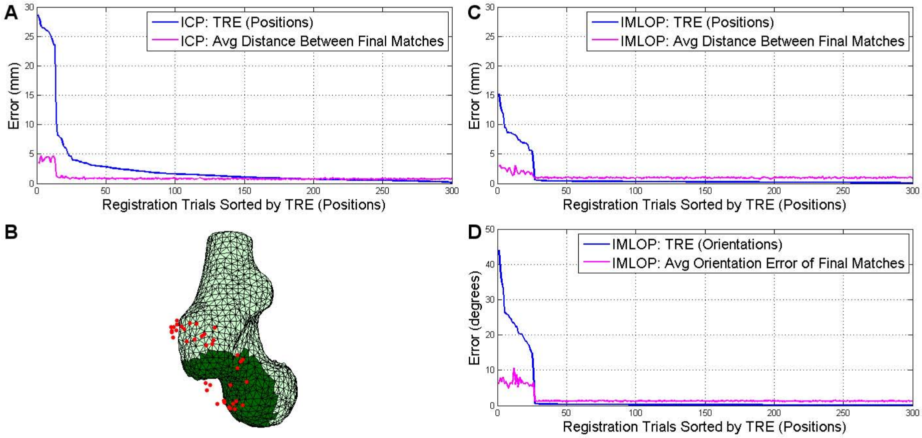

Fig. 1. Femur mesh and an example misaligned source shape point cloud (B); TRE and average final match error for the test case involving 75 samples and 1 mm (degree) noise level. Trials are sorted by positional TRE for both ICP (A) and IMLOP (C,D).

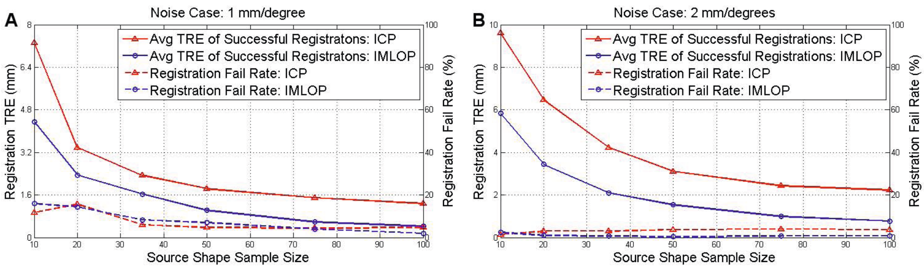

Fig. 2. Average TRE of successful registrations and registration failure rates across all sample sizes for noise level 1 mm (degree) (A) and 2 mm (degrees) (B)

4 Discussion and Conclusion

We have presented a novel algorithm for 3D-3D shape registration that extends the popular ICP method using a probabilistic framework to incorporate both position and orientation information. In addition, we have devised an efficient search strategy for computing the most probable oriented point correspondences and incorporated an efficient, closed-form solution for computing the optimal alignment of corresponding oriented points in the registration phase. As shown in the experimental results, the proposed IMLOP algorithm achieves significantly higher accuracy and has the advantage of a more robust mechanism for detecting registration failure by added criterion that is not avialable in ICP. Figure 1 shows that the registration errors of IMLOP decrease sharply when terminating close to the correct alignment, whereas ICP errors remain widely distributed. We have observed that IMLOP computes a solution in less than half the runtime of a similarly programmed ICP implementation, which demonstrates both a rate of convergence advantage for our method as well as the efficiency of our search strategy for computing the most probable oriented point correspondences.

Acknowledgments. Funded by the National Science Foundation Graduate Research Fellowship Program and Johns Hopkins University internal funds.

References

1. Rusinkiewicz, S., Levoy, M.: Efficient variants of the icp algorithm. In: Proceedings of the Third International Conference on 3-D Digital Imaging and Modeling, pp. 145–152 (2001)
2. Besl, P., McKay, N.D.: A method for registration of 3-d shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence 14, 239–256 (1992)
3. Arun, K., Huang, T.S., Blostein, S.D.: Least-squares fitting of two 3-d point sets. IEEE Transactions on Pattern Analysis and Machine Intelligence 9, 698–700 (1987)
4. Pulli, K.: Multiview registration for large data sets. In: Proceedings of the Second International Conference on 3-D Digital Imaging and Modeling, pp. 160–168 (1999)
5. Lara, C., Romero, L., Calderón, F.: A robust iterative closest point algorithm with augmented features. In: Gelbukh, A., Morales, E.F. (eds.) MICAI 2008. LNCS (LNAI), vol. 5317, pp. 605–614. Springer, Heidelberg (2008)
6. Kang, X., Armand, M., Otake, Y., Yau, W.P., Cheung, P., Hu, Y., Taylor, R.: Robustness and accuracy of feature-based single image 2D-3D registration without correspondences for image-guided intervention. IEEE Transactions on Biomedical Engineering 61, 149–161 (2014)
7. Liu, X., Cevikalp, H., Fitzpatrick, J.M.: Marker orientation in fiducial registration. In: Medical Imaging 2003, International Society for Optics and Photonics, pp. 1176–1185 (2003)
8. Banerjee, A., Dhillon, I.S., Ghosh, J., Sra, S., Ridgeway, G.: Clustering on the unit hypersphere using von mises-fisher distributions. Journal of Machine Learning Research 6 (2005)
9. Mardia, K., Jupp, P.: Directional Statistics. Wiley Series in Probability and Statistics. Wiley (2009)