

Fast and Robust Registration of 3D Surfaces Using Low Curvature Patches

Van-Duc Nguyen
GE Research & Development
Schenectady, NY 12301

Victor Nzomigni
CMA Consulting
Schenectady, NY 12301

Charles V. Stewart
Rensselaer Polytechnic Institute
Troy, NY 12180

Abstract

This paper describes a new model to range data registration algorithm, specifically designed for accuracy, speed, and robustness. Like many recent registration techniques, our Robust-Closest-Patch algorithm (RCP) iteratively matches model patches to data surfaces based on the current pose and then re-estimates pose based on these matches. RCP has several novel features: 1) on-line registration is driven by low curvature patches computed from the model off-line; 2) an approximate normal distance between a patch and a surface is used, avoiding the need to estimate local surface normal and curvature from noisy data; 3) pose is solved exactly by a linear system in six parameters, using a symmetric formulation of the rotation constraint; 4) robustness is ensured using an M-estimator that estimates both the rigid pose parameters and the error standard deviation. Results are shown using models and range data from turbine blade inspection.

1 Introduction

Automated inspection must be done both quickly and accurately. We are currently developing a blade inspection system which satisfies these two apparently conflicting requirements by: 1) measuring blade surfaces using a fast but low accuracy range sensor returning very dense surface points, 2) registering 100,000's of these points to get an accurate global registration, 3) subtracting these dense points to the registered CAD surfaces to find shape deformations. The 3D rigid transformation from the CAD model to the actual blade is typically found in less than 5 seconds on 400Mhz PC, and accurate to a standard deviation of 1/100 the sensor noise.¹

¹We are part of the Consortium for Non-Contact Gauging, whose goal is to rapidly and affordably measure aerospace and automotive parts 100 times faster than today's coordinate measuring machines (CMMs) and 10 times more accurately than conventional gauging.

Funded with a NIST ATP grant, the consortium's Rapid Agile Metrology for Manufacturing (RAMM) project incorporates advanced range sensor and software with CMM technology, to develop a shop floor system that can inspect loosely fixtured parts by scanning the surfaces and then aligning and comparing them to 3D CAD model for both in-process and final inspections.

The registration problem has received considerable attention, especially over the past five years, because it must be solved in numerous application domains, including medicine [9], object recognition [5, 7], robotic navigation [16], inspection [18], and automatic model construction [2, 4, 11, 17]. In spite of this work, a new solution, building on existing techniques, is needed to meet the demands of our application. Our solution must be accurate, extremely efficient, and robust to gross errors, often as many as 10-15% of the points.

The novel features of our algorithm can be seen by comparing it to related work, most notably Besl and McKay's Iterated Closest Point (ICP) algorithm [3] for registering 3D data with a model and Chen and Medioni's algorithm [4] for registering multiple range images. Several variations and improvements [2, 6, 16, 17, 22] on these algorithms have been shown, most complete is Stewart et al's covariance-based framework and generalization [20]. Briefly, given an initial rigid pose, all of these algorithms match each data point to its nearest point on the model surface or in the second data set, compute the 3D rigid transformation that best aligns these matches, and repeats these two steps using the most recently estimated pose until convergence. Our algorithm works in a similar framework, but has the following novel features:

1. Low curvature patches are precomputed off-line from the CAD model. These patches drive the on-line registration process from model to data, minimizing the required computation and allowing enough patches to be used to obtain the required accuracy.
2. An approximate normal distance between model patches and data surfaces is used, avoiding the need to estimate local surface normal and curvature from noisy data. A theoretical framework is developed to

Members of the consortium include: Brown & Sharpe, CMM manufacturer and system integrator; GE Corporate R&D, developer of software for part registration, surface reconstruction, and shape deformation analysis; Intelligent Automation Systems, developer of the 4DI coarse range sensor; other users such as: GE Aircraft Engines for airfoils, Caterpillar for crank shafts, Eaton for gears.




Figure 1: Example model and range data.

prove the approximation and to contrast it with previous distance constraints.

3. Pose is solved by singular value decomposition of a linear system in six parameters. The simplicity comes from a linear and symmetric formulation of the rotation constraint rather than quaternions or orthonormal matrices.
4. Robustness to gross errors in the data is ensured with an M-estimator that estimates both the rigid pose parameters and the error standard deviation.

2 Robust Closest Patch Algorithm

This section summarizes our Robust Closest Patch (RCP) algorithm and presents experimental results. Novel and theoretically important aspects of the normal distance constraint and of the pose solution are discussed in Sections 3 and 4 respectively.

2.1 Sketch of Algorithm

Like most other recent algorithms for registering 3D data sets or for registering 3D data with a model, RCP iterates two steps until convergence: it matches model and data based on the current pose estimate, and then refines the pose estimate based on these matches.

Much of the computation is done off-line by pre-computing low curvature patches. Given the approximate view point, the model surface can be digitized at regular grid points on an image plane, local curvature checked, and low curvature patches retained, giving a set of regularly spaced patches. Each patch P_i is represented by its center position \mathbf{p}_i and its outward normal \mathbf{n}_i . Figure 1a shows a shaded range image of the CAD model, seen from an a priori view point. Normal and occluding patches are shown by their surface normals. Figure 1b shows a range image, interpolated from triangulated range data points.

The on-line computation for a given range data set works as follows:

A.1 RCP translates the model patches to align the model's center of mass with that of the data. This is sufficiently robust because the object can be isolated in the data. Initial rotation is given by the assumed view point.

A.2 RCP repeats the following until convergence:

A.2.1 For each patch P_i , it finds the matching location \mathbf{q}_i by moving a matched filter, sized to P_i , along the line l_i , through \mathbf{p}_i and parallel to \mathbf{n}_i , searching for the nearest significant response from current location \mathbf{p}_i (Figure 6). This estimates the piercing point of l_i with the implicit surface from which the data are measured without the expensive and noise sensitive process of estimating surface parameters from the data [4]. As shown in Section 3, match point \mathbf{q}_i is a good approximation of the ideal match for \mathbf{p}_i if the patch P_i has low curvature.

A.2.2 RCP estimates the incremental rigid transformation, $\Delta T = (R, \mathbf{t})$, by minimizing the sum of the squared, approximate normal distances (8) between model patches and data surfaces:

$$\min_{\omega, \mathbf{t}_w} \sum_i [((I - W/2)\mathbf{q}_i - (I + W/2)\mathbf{p}_i - \mathbf{t}_w)^T \mathbf{n}_i]^2. \quad (1)$$

In this norm, error is calculated by rotating the data and model locations, \mathbf{q}_i and \mathbf{p}_i , each halfway toward the other, then taking the dot product of their difference along the model normal \mathbf{n}_i . The least-squares pose minimization and its robust version are discussed in Section 4.

Instead of using a quaternion or a small angle approximation, rotation R is represented by the skew symmetric matrix W , which describes the cross-product between a 3-element rotation vector $\omega = 2 \tan \frac{\theta}{2} \mathbf{u}$ and either \mathbf{p}_i or \mathbf{q}_i . θ is the angle and \mathbf{u} is the unit vector of the axis of rotation. The rotation vector ω and translation \mathbf{t}_w are solved directly from (1) with singular value decomposition, and then ΔT is:

$$\begin{aligned} R &= (I - W/2)^{-1}(I + W/2) \\ \mathbf{t} &= (I + W/2)\mathbf{t}_w \end{aligned} \quad (2)$$

A.2.3 RCP updates the positions and normals of the model patches and accumulates the incremental transformation into the current pose:

$$\begin{aligned} \mathbf{p}_i &= R\mathbf{p}_i + \mathbf{t} \\ \mathbf{n}_i &= R\mathbf{n}_i \\ T_{k+1} &= \Delta T \times T_k \end{aligned} \quad (3)$$

This entire registration process is embedded in a multi-resolution framework for global convergence and speed,




Figure 2: Absolute normal distance before and after registration.

using a small number of large patches at the coarsest resolution to eliminate the largest registration errors and lock into the global solution, and a large number of small patches at the finest resolution for precise registration. RCP has approximately linear time complexity in the number of model patches.

2.2 Experimental Results

RCP has been tested on 300 range data sets from 25 actual blades and 4 blade models with size varying from 2 to 6 inches. To illustrate, Figure 1 shows the CAD model and range data of a 3 inch blade. The “dovetail”, appearing below the platform in (a), is missing from (b). Also, the blade in (b), which is “in process” and has not been trimmed to its actual size, is longer than the model. Figure 2 shows the model and actual data sets over top of each other, then the absolute of the normal distance computed, before and after step 2 in RCP algorithm. Darker sections show larger deviations. During the iterations, RCP translated the model down and rotated it about the vertical axis, aligning the model and data platforms as well as their blade surfaces. RCP effectively ignored the extra length on the actual blade and treated missing sections of the model’s dovetail as outlier patches. Errors in alignment appear on the high-curvature blending areas where the airfoil and platform meet.

Figure 3a shows original range data, with missing data shown as white. Figure 3b shows the location of triangulated points whose residual error is greater than 3σ . Note how they form patches of different sizes, and how they occur close to missing data, or have larger deviations near high curvature areas. These outlier points can only be detected through solving for the robust pose of the object. Their effect on the root mean square error is shown in Figure 4, with the first increase caused by random noise, and




Figure 3: Noisy range data, and outliers with $d(P, P') > 3\sigma$.

the second increase by outliers.

Several quantitative measures of RCP’s performance can be given:

- Figure 5 shows the convergence of the robust RMS error, in mils, between model patches and data surfaces, as a function of the number of iterations of the registration process. The initial rotation was 45 degrees about the vertical axis. RCP typically converges in 4-5 iterations at each level of resolution.
- RCP converges quadratically if the robust weights (16) are fixed, typically if registration is within 1 pixel from its local minimum.
- RCP has correctly registered models with data containing up to 15% outliers. It should be able to handle higher percentages, but we have not seen these in our data.
- Tests of RCP have shown it correcting rotation errors of up to 60 degrees around each axis, separately. Of course, this large domain of convergence depends in part on the simple structure of the blade surfaces [3], and on the overconstrained minimization with both regular patches from visible surfaces and occluding patches from outlines.

3 Constraints from Patches

To derive the least squares norm (1) used in RCP, this section first introduces an ideal, symmetric distance measure (5) between model patches and data surfaces and shows how its first-order approximation (8) used in RCP eliminates the need to estimate local surface normal and curvature from noisy data. It then demonstrates quantitatively why low curvature patches are superior to high curvature landmarks such as corners and ridges, in terms of




Figure 4: Effect of deformations, noise and outliers on RMS error.




Figure 5: Convergence for 45° rotation about vertical axis.

representation completeness, registration accuracy, noise tolerance, and surface coverage. Finally, low curvature patches can also be used to capture the occluding surfaces along the object's outline.

3.1 Distance Along Average Normal

Let patch P be on the model surface, with mid point \mathbf{p} and normal \mathbf{n} defined a priori (Figure 6). Let patch P' be inferred from the range data as part of a larger planar surface with normal \mathbf{n}' . Since the local patch P' has no detectable boundaries or landmarks, its mid point \mathbf{p}' can not be found by local feature detection and matching. Instead, the location \mathbf{p}' corresponding to point \mathbf{p} is constructed to be the projection of \mathbf{p} onto the infinite plane supporting patch P' , along the average normal $\bar{\mathbf{n}}$:




Figure 6: Deviation between planar patch and surface.




Figure 7: Tangential sliding of patches leads to faster convergence.

$$\begin{aligned}\bar{\mathbf{n}} &= \frac{\mathbf{n} + \mathbf{n}'}{\|\mathbf{n} + \mathbf{n}'\|} = \frac{\mathbf{n} + \mathbf{n}'}{2 \cos(\alpha/2)} \\ \alpha &= \arcsin(\|\mathbf{n} \wedge \mathbf{n}'\|)\end{aligned}\quad (4)$$

The average of the two normals is meaningful if $\mathbf{n}^T \mathbf{n}' > 0$, which formalizes our assumption of rough alignment. Instead of constructing \mathbf{p}' based on the average normal $\bar{\mathbf{n}}$, intersection points \mathbf{q} and \mathbf{q}' can be constructed by projecting \mathbf{p} along the normal \mathbf{n} [4] and \mathbf{n}' respectively. All three are related, and so the distance measure between a model patch P and a data surface containing patch P' can be constructed as:

$$\begin{aligned}d(P, P') &= (\mathbf{p}' - \mathbf{p})^T \bar{\mathbf{n}} \\ &= ((\mathbf{q} - \mathbf{p})^T \mathbf{n}) (\cos \alpha / \cos(\alpha/2)) \\ &= ((\mathbf{q}' - \mathbf{p})^T \mathbf{n}') / \cos(\alpha/2)\end{aligned}\quad (5)$$

Note that the patch P' is free to rotate and translate tangentially along its surface. Minimizing the normal distance between point and surface leads to much faster convergence [4] to the final pose than minimizing the Euclidean distance between model and matched points [3]. The translation and rotation of a vertical rectangle in Figure 7 requires only one instead of many registration iterations (matching and pose refinement) if Euclidean point distances are used: the convergence in the translation example is at best quadratic. In both translation or rotation examples, Euclidean point distances along the vertical edges of the rectangle restrict tangential sliding of these vertical patches, leading to not only slow convergence but also multiple weak minima near the global minimum error.

3.2 Approximate Normal Distance

Let r be the local radius of curvature at the patches P and P' (Figure 8). The points \mathbf{q} and \mathbf{q}' are now on the curved patch P' , and are found from the mid point \mathbf{p} of patch P , along the normals \mathbf{n} and \mathbf{n}' respectively. The distance measure between a curved patch and a curved surface, is defined as:

$$\begin{aligned}d(P, P') &= (\mathbf{p}' - \mathbf{p})^T \bar{\mathbf{n}} \\ &\approx ((\mathbf{q} - \mathbf{p})^T \mathbf{n} + r\beta^2/8) (\cos \alpha / \cos(\alpha/2)) \\ &\approx ((\mathbf{q}' - \mathbf{p})^T \mathbf{n}' + r\beta^2/8) / \cos(\alpha/2) \\ \beta &\approx \alpha d(P, P') / r\end{aligned}\quad (6)$$




Figure 8: Deviation between curved patch and surface.

The simplicity and symmetry in the above approximations depends on the assumption that the local radius of curvature r is greater than the absolute normal distance $|d(P, P')|$. In particular, the curved arc between points \mathbf{q} and \mathbf{q}' makes an angle α at the mid point \mathbf{p} , and an angle β at the center of curvature of the patch P' . Using Taylor expansions and eliminating β , local curvature is shown to only add second order terms in α :

$$\begin{aligned} d(P, P') &\approx (\mathbf{q} - \mathbf{p})^T \mathbf{n} (1 + (d(P, P')/r - 3) \alpha^2/8) \\ &\approx (\mathbf{q}' - \mathbf{p})^T \mathbf{n}' (1 + (d(P, P')/r + 1) \alpha^2/8) \end{aligned} \quad (7)$$

When $\alpha \approx 0$ and $r > |d(P, P')|$, the following approximate normal distance between low curvature patches and surfaces is used to avoid estimating local surface normal and curvature from range data:

$$d(P, P') \approx (\mathbf{q} - \mathbf{p})^T \mathbf{n} \quad (8)$$

3.3 Singular Patches

By definition a regular patch has a well defined surface normal. A ridge segment or a corner point are instances of surface singularity. The surface normal is non unique at exactly a ridge segment or a corner point. However, if the patches leading to the ridge (resp. corner) are planar, or have low enough curvature, then the normal distance constraints from these two (resp. three) patches are equivalent to a collinearity (resp. coincidence) constraint from the segment (resp. point) (Figure 9). For example, the Euclidean distance between two corner points P and Q is nothing but the aggregate of three normal distances along three perpendicular axes:

$$\begin{aligned} \delta(P, Q)^2 &= (\mathbf{q} - \mathbf{p})^T (\mathbf{q} - \mathbf{p}) \\ &= (q_x - p_x)^2 + (q_y - p_y)^2 + (q_z - p_z)^2 \quad (9) \\ &= \sum_{\mathbf{n} \in \{\mathbf{x}, \mathbf{y}, \mathbf{z}\}} ((\mathbf{q} - \mathbf{p})^T \mathbf{n})^2 \end{aligned}$$

The three normal distances along the three surface normals is a complete and much more accurate distance measure




Figure 9: Ridges and corners are singular patches.

than the Euclidean distance between corner locations, because the real normals are taken into account.

Furthermore, ideal corner points or ideal ridge segments do not exist in the real world (Figure 9). There is always a non zero radius of curvature caused by a blending surface to avoid concentration of stress which may lead to cracks. On a typical airfoil for aircraft engines and power turbines, this radius of curvature could be 10 to 1000 times the required accuracy of the surface. So using high curvature features such as corners and ridges would mean 10 to 1000 times less accurate registration. Or it would require full modeling of the blending surfaces, which are usually not well defined in the CAD model either.

3.4 Occluding Patches

The blade in Figure 1 is mainly composed of two large perpendicular surfaces: an airfoil and a platform. The platform is mostly planar. The airfoil is a low-curvature, ruled surface. This is an instance of shape singularity, in which the regular surfaces do not provide enough constraints to solve for all six pose parameters reliably. In particular, the ratio of maximum over minimum eigenvalue is 229.8. This ill-condition number decreases to 20.51 when occluding patches along the blade's outline are added. An advantage of the pre-compilation of the CAD model into patches is the ability to detect shape singularities, and take corrective actions ahead of actual registration. In general, both regular and occluding patches along with appropriate weights should fully constrain the global pose, except for global shape symmetries and reflections as in cylinders and ellipsoids.

4 Solving Pose with Normal Distance

Previous registration algorithms have minimized the sum of squares of approximate normal distances (8):

$$\min_{R^T, \mathbf{t}'} \sum_i [(R^T \mathbf{q}_i - \mathbf{p}_i - \mathbf{t}')^T \mathbf{n}_i]^2 \quad (10)$$

using a small angle approximation for R [2]. From the estimated inverse $(R^T, -\mathbf{t}')$, the transform $(R, R\mathbf{t}')$ is applied

to the model patches to align the CAD model to the range data.

Ideally, the pose should be estimated by robustly minimizing the sum of squares of symmetric normal distances (5):

$$\min_{R_{1/2}, \mathbf{t}_{1/2}} \sum_i [(R_{1/2}^T \mathbf{p}'_i - R_{1/2} \mathbf{p}_i - \mathbf{t}_{1/2})^T \bar{\mathbf{n}}_i]^2 \quad (11)$$

The error for match i is measured by rotating \mathbf{p}'_i and \mathbf{p}_i each halfway toward the other, then taking the dot product of the residual deviation with the average normal $\bar{\mathbf{n}}_i$. It is implicit that the two normals are also rotated halfway toward $\bar{\mathbf{n}}_i$. After estimating $R_{1/2}$ and $\mathbf{t}_{1/2}$, the transform applied to the model patches would be $R = R_{1/2} R_{1/2}$ and $\mathbf{t} = R_{1/2} \mathbf{t}_{1/2}$.

Unfortunately, neither (10) nor (11) can be manipulated as in [7, 13] to separate rotation and translation and allow rotation to be solved exactly in closed form using quaternions. This section proves how the symmetric error norm (11) is equivalent to the algebraic error norm (12) if matching is exact, then approximated to actual norm (13) used in RCP, by replacing the average normal $\bar{\mathbf{n}}_i$ with the model normal \mathbf{n}_i .

4.1 Algebraic Error Norm

The ideal minimization (11) can be rewritten using the linear pose constraint (21), yielding another least squares error norm,

$$\sum_i [(I - W/2) \mathbf{p}'_i - (I + W/2) \mathbf{p}_i - \mathbf{t}_w]^T \bar{\mathbf{n}}_i]^2, \quad (12)$$

which no longer depends on rotation matrices $R_{1/2}$. With this algebraic norm, the parameters ω and \mathbf{t}_w can be solved exactly with linear equations, from which the desired rotation R and translation \mathbf{t} can be calculated as in (2). And since there are no constraints on the parameters as in quaternions [7, 13], the linear system in the six pose parameters can be solved using a singular value decomposition.

The approximation $R_{1/2} \approx I + W/2$ introduces only radial errors, perpendicular to $\mathbf{p}' - \mathbf{p}$, and so generates no resulting moment along or against current rotation ω , Figure 10, i.e. has absolutely no effect on the least-squares-error solution. In other words, large rotations can be solved in one iteration if correct matches are assumed, and so the convergence is faster than solving with small angle approximation to linearize R .

Finally, using model normal \mathbf{n}_i rather than the average normal $\bar{\mathbf{n}}_i$, the approximate least squares error norm is

$$\sum_i [(I - W/2) \mathbf{q}_i - (I + W/2) \mathbf{p}_i - \mathbf{t}_w]^T \mathbf{n}_i]^2. \quad (13)$$

From (7), this approximation leads to over-estimating the normal distance by a factor of $(1 - (d(P, P')/r - 3) \alpha^2/8)$

which in turn leads to over-estimating the rotation vector ω . Despite (13) being an approximation to (12), there are still distinct advantages to using it. In our simulations, it is a good approximation even for large rotation angles of 90 degrees. For true rotations of less than 45 degrees, the overshoot in estimating rotation is slightly smaller than that obtained when solving (10) using a small angle approximation to R .

4.2 Robust Pose Estimation

The pose estimation procedure is made robust while still maintaining computational efficiency by using an M-estimator [10, 14]. Outliers, often in clusters, appear in the range data due to several causes: specular reflections from the shiny metallic surface, measurements that are not part of the blade, and regions of the actual blade surface which do not appear in the model. These can not be trimmed by restricting data to a hypothetical work volume. Preprocessing and filtering the range data is computationally inefficient, ineffective (when outliers are clustered), and unnecessary because robustness is easily built into the pose calculation. Finally, the recent use of high breakdown point estimators such as least median of squares (LMS) [19] in the pose calculation [17] is too expensive and unnecessary in our application to obtain a high degree of robustness (see [15] for related discussion).

To use an M-estimator effectively we must calculate a scale estimate σ in addition to the pose parameters. We discuss pose estimation first, assuming fixed σ . The M-estimate of pose is

$$\operatorname{argmin}_{\omega, \mathbf{t}_w} \sum_i \rho(r_i(\omega, \mathbf{t}_w)/\sigma). \quad (14)$$

where the residual error r_i describes the approximate normal distance as in (13):

$$r_i(\omega, \mathbf{t}_w) = ((I - W/2) \mathbf{q}_i - (I + W/2) \mathbf{p}_i - \mathbf{t}_w)^T \mathbf{n}_i \quad (15)$$

This is solved using iteratively reweighted least-squares (IRLS) [12], with weighting function $w(r/\sigma) = \rho'(r/\sigma)/(r/\sigma)$. Our current implementation uses the ‘‘Tukey biweight’’ [1]:

$$\rho(r/\sigma) = \begin{cases} \frac{B^2}{6} [1 - (1 - (r/\sigma)^2)^3] & |r/\sigma| \leq B \\ \frac{B^2}{6} & |r/\sigma| > B \end{cases} \quad (16)$$

whose tuning constant B is typically around 4.5 to ensure consistency at Gaussian errors [12]. The resulting weight function is 0 for residuals larger than $B\sigma$.

Turning now to scale estimation, it is crucial to recalculate σ during each iteration of the registration process. For, the normal residuals r_i reflect both pose error and sensor noise, with pose error dominating initially and noise

dominating at the final pose. Overall, the scale of the errors can decrease by two orders of magnitude. Thus, scale estimated from the initial pose error will be unable to eliminate outliers as registration converges, and setting σ initially based on the known sensor error will eliminate most matches as outliers. To avoid this problem, scale is estimated twice during each iteration of the registration process. After matching and before running IRLS, σ is estimated from the current pose and then used in IRLS. After IRLS, σ is re-calculated from the estimated pose and used to limit the search for matches in the next iteration of registration, providing additional robustness and efficiency. For any given set of matches and pose estimate, scale is calculated using the median absolute deviation [12]:

$$\sigma = 1.483 \operatorname{median}_i \{|r_i(\boldsymbol{\omega}, \mathbf{t}_w)|\}, \quad (17)$$

where the constant multiplier ensures consistency at Gaussian residuals.

5 Appendix

5.1 Linear Solution of Pose

Let $\mathbf{p}' = R\mathbf{p}$, assuming pure rotation by an angle θ around an axis described by the unit vector \mathbf{u} . Figure 10 shows the following formula of Rodrigues [8, p. 72] [13]:

$$\mathbf{p}' = (\mathbf{p}^T \mathbf{u}) \mathbf{u} + (\mathbf{p} - (\mathbf{p}^T \mathbf{u}) \mathbf{u}) \cos \theta + (\mathbf{u} \wedge \mathbf{p}) \sin \theta \quad (18)$$

Rotating \mathbf{p}' back to \mathbf{p} requires just reversing the sign of θ . Using this to form an equation similar to (18), subtracting the two equations, then using the fact that $\mathbf{p}' - \mathbf{p}$ is normal to \mathbf{u} and rearranging slightly yields the direct solution of rotation:

$$\begin{aligned} (\mathbf{p}' - \mathbf{p}) &= \boldsymbol{\omega} \wedge (\mathbf{p} + \mathbf{p}')/2 \\ \boldsymbol{\omega} &= 2 \tan \frac{\theta}{2} \mathbf{u} \end{aligned} \quad (19)$$

The axis of rotation \mathbf{u} and the angle of rotation θ are all described in the 3-component vector $\boldsymbol{\omega}$ with no additional constraints on it. The rotation vector $\boldsymbol{\omega}$ has a singularity at $\theta = \pi$, but no singularity at $\theta = 0$, unlike a quaternion [13]. The singularity at $\theta = \pi$ is unavoidable since all rotations can not be represented using just three parameters [21].

Now, when rotation and translation are both involved, $\mathbf{p}' = R\mathbf{p} + \mathbf{t}$. Substituting $\mathbf{p}' - \mathbf{t}$ in place of \mathbf{p}' in (19) yields the direct solution of the full pose:

$$\begin{aligned} (\mathbf{p}' - \mathbf{p}) &= \boldsymbol{\omega} \wedge (\mathbf{p} + \mathbf{p}')/2 + \mathbf{t}_w \\ \mathbf{t}_w &= \mathbf{t} - \frac{1}{2} \boldsymbol{\omega} \wedge \mathbf{t}. \end{aligned} \quad (20)$$

5.2 Relation to Orthonormal Matrices

Let W be the skew symmetric matrix such that $\boldsymbol{\omega} \wedge \mathbf{v} = W\mathbf{v}$ for all vectors \mathbf{v} in \mathfrak{R}^3 . Then (20) can be rewritten as

$$\begin{aligned} (I - W/2)\mathbf{p}' &= (I + W/2)\mathbf{p} + (I - W/2)\mathbf{t} \\ \mathbf{p}' &= (I - W/2)^{-1}(I + W/2)\mathbf{p} + \mathbf{t}_w \end{aligned} \quad (21)$$




Figure 10: Illustrating Rodrigues' formula.

$(I - W/2)$ is always invertible and $(I - W/2)^{-1}(I + W/2)$ is an orthonormal matrix. So, the relationship between the rotation vector $\boldsymbol{\omega}$ and its equivalent orthonormal matrix R is

$$R = (I - W/2)^{-1}(I + W/2) \quad (22)$$

It can be shown that $I - W/2$ and $I + W/2$ are very close to orthonormal for small angle of rotation θ , and so $R_{1/2}^T \approx I - W/2$ and $R_{1/2} \approx I + W/2$.

References

- [1] A. E. Beaton and J. W. Tukey. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*, 16:147–185, 1974.
- [2] R. Bergevin, M. Soucy, H. Gagnon, and D. Laurendeau. Towards a general multiview registration technique. *IEEE Trans. on PAMI*, 18(5):540–547, 1996.
- [3] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Trans. on PAMI*, 14(2):239–256, 1992.
- [4] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.
- [5] C. Chua and R. Jarvis. 3d free-form surface registration and object recognition. *Int. J. of Computer Vision*, 17(1):77–99, 1996.
- [6] C. Dorai, J. Weng, and A. Jain. Optimal registration of multiple range views. In *Proc. ICPR*, pages A:569–571, 1994.
- [7] O. Faugeras and M. Hebert. The representation, recognition, and locating of 3-d objects. *Int. J. of Robotics Research*, 5(3):27–52, 1986.
- [8] I. D. Faux and M. J. Pratt. *Computational Geometry for Design and Manufacture*. Ellis Horwood Ltd., 1979.
- [9] W. Grimson, T. Lozano-Perez, W. Wells, G. Ettinger, and S. White. An automatic registration method for frameless stereotaxy, image, guided surgery and enhanced reality visualization. In *Proc. CVPR*, pages 430–436, 1994.
- [10] F. R. Hampel, P. J. Rousseeuw, E. Ronchetti, and W. A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. John Wiley & Sons, 1986.

- [11] K. Higuchi, M. Hebert, and K. Ikeuchi. Building 3-d models from unregistered range images. *Graphical Models and Image Processing*, 57(4):315–333, 1995.
- [12] P. W. Holland and R. E. Welsch. Robust regression using iteratively reweighted least-squares. *Commun. Statist.-Theor. Meth.*, A6:813–827, 1977.
- [13] B. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. Amer.*, 4(4):629–642, 1987.
- [14] P. J. Huber. *Robust Statistics*. John Wiley & Sons, 1981.
- [15] R. Kumar and A. R. Hanson. Robust methods for estimating pose and a sensitivity analysis. *CVGIP: Image Understanding*, 1994.
- [16] F. Lu and E. Milios. Robot pose estimation in unknown environments by matching 2d range scans. In *Proc. CVPR*, pages 935–938, 1994.
- [17] T. Masuda and N. Yokoya. A robust method for registration and segmentation of multiple range images. *Computer Vision and Image Understanding*, 61(3):295–307, 1995.
- [18] C.-H. Menq, H.-T. Yau, and G.-Y. Lai. Automated precision measurement of surface profile in cad-directed inspection. *IEEE Trans. on Robotics and Automation*, 8(2):268–278, 1992.
- [19] P. J. Rousseeuw. Least median of squares regression. *J. Amer. Stat. Assoc.*, 79:871–880, 1984.
- [20] C. V. Stewart, K. Bubna, V.-D. Nguyen, and V. Nzomigni. Covariance based registration. In *Proc. CVPR*, will be published in 2000.
- [21] J. Stuelpnagel. On the parameterization of the three-dimensional rotation group. *SIAM Review*, 6(4):422–430, 1964.
- [22] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Int. J. of Computer Vision*, 13(2):119–152, 1994.