

Verification of Multi-View Point-Cloud Registration for Spherical Harmonic Cross-Correlation

Robert L. Larkins
rll6@students.waikato.ac.nz

Michael J. Cree
cree@waikato.ac.nz

Adrian A. Dorrington
adrian@waikato.ac.nz

School of Engineering
University of Waikato
Private Bag 3105, Hamilton 3240, New Zealand

ABSTRACT

Spherical harmonic cross-correlation is a robust registration algorithm that brings two point-clouds of the same scene into coarse rotational alignment. The found rotation however may not give the desired alignment, as misalignments can occur if there is not enough overlap between point-clouds, or if they contain a form of symmetry. We propose a verification method whose purpose is to determine if registration has failed for a priori unknown registration. The rotational transformation between multiple clouds must satisfy internal consistency, namely multiple rotational transformations are transitive. The rotation verification is performed using triplets of images, which are cross-referenced with each other to classify rotations individually. Testing is performed on a dataset of a priori known registrations. It is found that when the number of images or the percentage of correct rotations is increased, the number of correct rotation classifications improves. Even when tested with only four images and a correct rotation percentage of 17%, the rotation verification is still considered a viable method for classifying rotations. Spherical harmonic cross-correlation is benefited by rotation verification as it provides an additional approach for checking whether found rotations are correct.

Categories and Subject Descriptors

I.4.3 [Image Processing and Computer Vision]: Registration

Keywords

3D Imaging, Registration, Point Cloud, Spherical Harmonics, Rotation Verification

1. INTRODUCTION

Range imaging systems [2] typically image an object or scene from a single pose, therefore, only part of the scene is acquired. More of the scene is revealed by shifting the

system and acquiring subsequent images. As these images are three-dimensional (3D) in nature, they are expressed as point-clouds. Registration is the process of aligning these images in the same coordinate space, building a full 3D representation. However, if there is no method of tracking how the imaging system or object has moved, then the transformation between image captures is a priori unknown. In this situation, the overlap between point-clouds is used as the basis for performing registration. A variety of algorithms exist that use overlap to perform registration, such as iterative closest point (ICP), random sample consensus or principal component analysis [10]. Spherical harmonic cross-correlation (SHCC) [5] is an algorithm that achieves coarse rotational alignment and is robust to noise. The ascertained rotation, however, is not guaranteed to give the desired alignment, as misalignments can occur if the overlap is insufficient, or the data contains symmetry.

Two methods that can determine whether the SHCC alignment is correct is to threshold either the resulting error between the two sampled surfaces [6], or the maximum impulse value from the Fourier phase correlation when performing translational alignment. Their disadvantage is that they must be performed separately from SHCC, whereas our focus is to verify the calculated rotations using the rotations themselves.

The transitive property of rotations, namely that two rotations in sequence are the same as one complete rotation, can be used to check the consistency of the SHCC registrations. Thus if we have three images of an object, the three possible image pairings each produce a rotation; each rotation must be expressible as a combination of the other two, otherwise it may be concluded that one or more of the rotations are incorrect. We propose applying verification via transitivity to multiple groupings of images from the same set to determine the incorrect rotational registrations without a priori knowledge of the transformations between poses.

There appears to be a dearth of literature that applies transitivity to point-cloud registration. For registering two-dimensional medical images, only the transformations that maps one image to another and vice versa are considered [3], as this ensures that these transformations are consistent. This approach is expanded on by Geng et al. [4], who iteratively find the registration of three manifolds using the closest point projection to find correspondences, and are improved by minimising the error in the transitive inverse. Transformational transitivity has been utilised for a multi-view extension to the iterative closest point algorithm [9], identifying transformations that link non-adjacent images,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IVCNZ '12, November 26 - 28 2012, Dunedin, New Zealand
Copyright 2012 ACM 978-1-4503-1473-2/12/11 ...\$15.00.

helping minimise the alignment error of motion averaging.

The work presented investigates whether the rotations found from the SHCC can be verified as correct or incorrect by testing for transitivity. This is performed using triplets of images and analyses how the size of the image set and the ability of the SHCC impacts the classification accuracy of the verification.

2. REGISTRATION PRELIMINARIES

The acquired point-clouds are given in Cartesian coordinates, in which the coordinate space is attached to the imaging system, such that the z-axis is perpendicular to the imaging sensor and points into the scene. SHCC attempts to identify a rotation between two overlapping point-clouds so that one point-cloud can be rotated into the coordinate space of the other (or into world coordinates). This rotation can be interchangeably referred to as the rotation that aligns the two point-clouds, or as the rotation separating the two poses of the imaging system. The rotational alignment is then complemented by the translational alignment, fully registering the two point-clouds.

2.1 Spherical Harmonic Cross-Correlation

The SHCC uses the surface normals of two overlapping point-clouds to find an alignment rotation; these normals are ascertainable when the relationship between points is known, making SHCC applicable to alternative surface representations, such as meshes. The normals are converted into a set of spherical harmonic functions using associated Legendre polynomials, allowing the correlation to be performed in the frequency domain. Let P_l^m be the associated Legendre polynomial of degree l with $l \in 0, \dots, L$, and of order m with $m \in -l, \dots, l$, where L is the maximum spherical harmonic degree. The spherical harmonics of a normal $N_{[\theta, \phi]}$, are given by

$$Y_l^m(N) = \begin{cases} \sqrt{\frac{(2l+1)(l-m)!}{4\pi(l+m)!}} P_l^m(\cos N_\theta) e^{imN_\phi} & m \geq 0, \\ (-1)^m Y_l^{|m|}(N) & m < 0, \end{cases} \quad (1)$$

where \bar{x} is the complex conjugate of x . If the set of n_N normals from the first acquisition is labelled f , the spherical harmonic representation of f is found as

$$\hat{f}_l^m = \sum_{j=0}^{n_N-1} Y_l^m(N_j), \quad (2)$$

with the analogous operation applied to the normals g of the second point-cloud to give \hat{g} . The Fourier transform of the three-dimensional rotation correlation matrix, \hat{C}_R , between \hat{f} and \hat{g} is given by

$$\hat{C}_R(m, n, k) = \sum_{l=\max(|m|, |n|, |k|)}^L \overline{\hat{f}_l^m} \hat{g}_l^n d_{mk}^l d_{kn}^l, \quad (3)$$

for $m, n, k \in -L, \dots, L$, where the function d [8] is

$$d_{mn}^l = 2^{-l} \sum_{t=\max(0, n-m)}^{\min(l+n, l-m)} (-1)^t \frac{\sqrt{(l+n)!(l-n)!(l+m)!(l-m)!}}{(l+n-t)!(l-m-t)!(t+m-n)!t!}. \quad (4)$$

The final step is to apply the inverse Fourier transform to \hat{C}_R to obtain the correlation matrix, C_R , with the greatest impulse response of C_R indicating the rotation that maximises the overlap between the two sets of normals. Taking a, b and c as the indices of the greatest value in C_R , the Euler angles α, β and γ are found by

$$\alpha = \text{mod}(sa + \pi/2, 2\pi), \quad (5)$$

$$\beta = \text{mod}(sb + \pi, 2\pi), \quad (6)$$

$$\gamma = \text{mod}(sc + \pi/2, 2\pi), \quad (7)$$

where $s = 2\pi/(2L+1)$, the size of each rotational division. Applying these angles to a ZYZ rotation matrix will generally bring the two point-clouds into the correct rotational alignment.

2.2 Translational Alignment

Once the two overlapping point-clouds are in rotational alignment, their translational alignment can be calculated using a three-dimensional Fourier phase correlation. As the rotation verification is performed without needing to know the translational alignment, it is not described with detail in this paper. Larkins et al. [8] provides an outline of the mathematics that give translational alignment.

3. ROTATION VERIFICATION

The rotations that separate the poses of an imaging system can be established as correct if they are consistent with each other. A set of n_A images produce

$$n_R = \frac{n_A(n_A - 1)}{2} \quad (8)$$

rotations that link these images together, where n_A is a non-negative integer.

If each image has a unique index value in the range of $0, \dots, n_A - 1$, then each rotation can be indexed using the two images that it links, allowing each rotation to be stored and accessed with ease. If two images are indexed a and b , and the total number of images is n , then the rotational index R_i is given as

$$R_i(a, b, n) = a \left(\frac{2n - a - 3}{2} \right) + b - 1, \quad (9)$$

$$0 \leq a < b < n.$$

The verification used here focuses on triplets of rotations, as three rotations are the minimum needed to establish rotational consistency. The number n_T of triplets produced from n_A images is given by

$$n_T = \frac{n_A(n_A - 1)(n_A - 2)}{6}. \quad (10)$$

Each triplet can be individually indexed via the three images from which it is comprised; this index is found as

$$T_i = \frac{j^3 + 3j^2(1 - n_A) + j(3n_A^2 - 6n_A + 2)}{6} + R_i(k - j - 1, l - j - 1, n_A - j - 1), \quad (11)$$

$$0 \leq j < k < l < n_A,$$

where j, k , and l are the indices of the three images. A triplets index facilitates the storing and retrieval of its classification, but depending on the implementation design, may not be necessary.

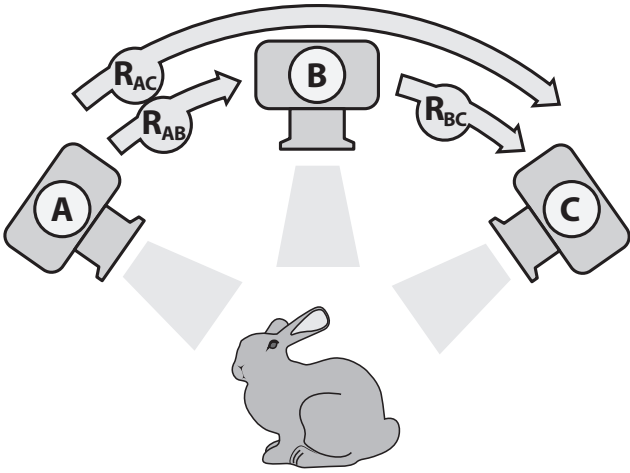


Figure 1: Three poses, A , B and C , of an imaging system are related to each other by the three rotations \mathbf{R}_{AB} , \mathbf{R}_{BC} and \mathbf{R}_{AC} . If the rotations found from registration are correct, then \mathbf{R}_{AC} is equivalent to \mathbf{R}_{AB} followed by \mathbf{R}_{BC} , as they both rotate the imaging system from A to C .

3.1 Triplet Verification

Each triplet is classified individually using its three rotations to be *aligned* if its rotations are consistent with each other, otherwise it is *misaligned*. Three point-clouds, A , B and C , and the three rotations, \mathbf{R}_{AB} , \mathbf{R}_{AC} and \mathbf{R}_{BC} , that register point-cloud A to B , and so on, then the three rotations can be said to be consistent if

$$\mathbf{R}_{AC} = \mathbf{R}_{BC}\mathbf{R}_{AB}. \quad (12)$$

This relationship is illustrated in Figure 1. However, there are two limitations that may result in an inconsistent triplet by the above definition, even if the rotational alignment is correct. The first is that if the rotations are represented as rotation matrices, then the rotations can be expressed in two ways, as a clockwise or counter-clockwise rotation about an arbitrary axis; combined with their noncommutative nature, the point-clouds must be labelled in a specific order for $\mathbf{R}_{BC}\mathbf{R}_{AB}$ to equal \mathbf{R}_{AC} , but as this order is unknown, there is no guarantee that inconsistency is proof of misalignment. The second is that rotations found via SHCC come from a discrete set, resulting in quantisation error, thus \mathbf{R}_{AC} and $\mathbf{R}_{BC}\mathbf{R}_{AB}$ may not precisely agree.

To arrive at a better definition of consistency, let

$$\mathbf{R}_S = \mathbf{R}_{AC}^T \mathbf{R}_{BC} \mathbf{R}_{AB}. \quad (13)$$

and

$$\theta = \cos^{-1} \left(\frac{\text{tr}(\mathbf{R}_S) - 1}{2} \right). \quad (14)$$

If

$$|\theta| \leq \epsilon \quad (15)$$

for some small margin-of-error ϵ , then the triplet comprising rotations \mathbf{R}_{AB} , \mathbf{R}_{AC} and \mathbf{R}_{BC} is *consistent*. If the triplet is inconsistent then at least one of the three rotations is incorrect; the converse is not necessarily true, as two or more incorrect rotations may coincidentally satisfy transitivity. A consistent triplet is classified as *aligned* and an inconsistent triplet is classified as *misaligned*. If the rotations of a triplet

align correctly, then the allowable error ϵ has an inverse relationship with the maximum spherical harmonic degree L , that is, as L increases, ϵ can be decreased, limiting false positives.

3.2 Classifying Individual Rotations

Triplet verification classifies triplets as aligned or misaligned. It remains to classify the individual rotations comprising a triplet. The classification of an individual rotation can be inferred from the $n_A - 2$ triplets that it contributes to. If all $n_A - 2$ triplets are classified as aligned, then it is highly likely the rotation has been correctly deduced, however, if the triplets are all misaligned, there is no guarantee that the rotation is incorrect as the misaligned triplets may be a result of other incorrect rotations. But as the rotation contributes only to misaligned triplets, it cannot be determined correct, thus is treated as incorrect. For all other situations a threshold is specified with the rotation being deemed correct if the number of aligned triplets it contributes to meets the threshold. Three thresholds of differing strictness are investigated herein, namely,

$$\tau_a = 1, \quad (16)$$

$$\tau_b = \left\lceil \frac{n_A - 1}{2} \right\rceil, \quad (17)$$

$$\tau_c = n_A - 2. \quad (18)$$

3.3 Verification Worst Case

Triplet verification fails when every triplet contains one or more incorrect rotations. If a rotation is incorrect, then all $n_A - 2$ triplets that it contributes to are misaligned. The minimum number of incorrect rotations that render all triplets misaligned is

$$q = \left\lceil (n_A - 1)^2 / 4 \right\rceil. \quad (19)$$

The total number of q rotation combinations is given by the binomial coefficient,

$$\binom{n_R}{q}. \quad (20)$$

Fortunately, complete misalignment only occurs for a subset of these combinations, of size

$$\binom{2h - 1}{h}, \quad (21)$$

where $h = \lfloor (n_A + 1) / 2 \rfloor$. An example of the incorrect rotation combinations for $n_A = 6$ are shown in Figure 2, in which the ten possible arrangements of the six incorrect rotations cause all triplets to be misaligned.

4. METHODOLOGY

Triplet verification is tested by a dataset that has been formed from 120 separate poses about the Dragon model¹, shown in Figure 3. Using uniform deviates $u, v \in [0, 1)$, each pose, A , given in spherical coordinates, is randomly placed about the model by

$$A_{\theta, \phi} = (2\pi u, \cos^{-1}(2v - 1)), \quad (22)$$

¹This model comes from the Stanford 3D scanning repository (<http://graphics.stanford.edu/data/3Dscanrep/>).




Figure 2: In the case of six images, fifteen rotations are formed. The shown combinations are three of the ten the worst case scenarios that occur when six rotation (solid lines) are incorrect. For six images, six incorrect rotations are the minimum required to cause complete triplet verification failure.




Figure 3: The Dragon model from the Stanford 3D scanning repository that is used for testing. The model is shown in its initial state before being broken into individual point-clouds.

with its optical axis directed towards the model. This viewing orientation is then further varied by randomly rotating it about the optical axis. Using Blender², the surface of the model visible to each imaging pose produces a point-cloud, as well as generating a normal for each point.

The SHCC is applied to every image pair, creating 7 140 rotations. As the acquiring pose of all images is known a priori, each rotation is deemed correct if the image is rotated to within 10° of its known pose, providing the ground truth classification; 10° is chosen as ICP, a fine registration algorithm, can register angular differences up to this size [7]. A maximum spherical harmonic degree of $L = 20$ is used, therefore, a margin-of-error of $\epsilon = 10^\circ$ is chosen to accommodate the combined quantisation error from two rotations. There are 1 182 correctly determined rotations and 5 958 incorrect rotations. The low number of correct rotations is due to many pairs of images having low overlap, as shown in Figure 4. Triplet verification is then used to classify each rotation that connects a set of images together. Triplet verification classification is then compared to the ground truth to establish the number of each possible outcome, which are given in Table 1. The true-positive fraction (TPF) and false-positive fraction (FPF) are calculated as

$$\text{TPF} = \text{TP}/(\text{TP} + \text{FN}), \quad (23)$$

$$\text{FPF} = \text{FP}/(\text{FP} + \text{TN}), \quad (24)$$

where TP is the number of true positives, TN the number

²Blender is an open-source 3D graphics editor (<http://www.blender.org/>). Version 2.62.0 r44136 is used here.

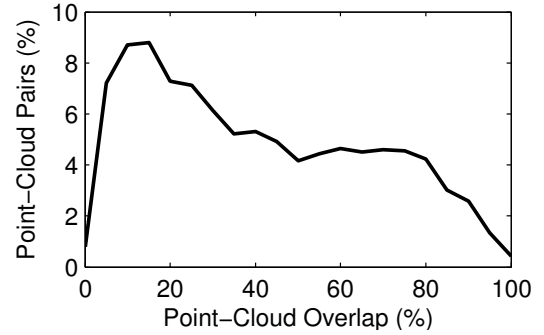


Figure 4: The percentage of point-cloud pairs that have a certain overlap percentage.

Verification Classification	Ground Truth Classification	
	Aligned	Misaligned
Aligned	True Positive	False Positive
Misaligned	False Negative	True Negative

Table 1: The four possible outcomes when verification classification is compared with the ground truth classification.

of true negatives, FP the number of false positives and FN the number of false negatives. Three factors that impact the verification classification are investigated in this paper; the number of images registered, the percentage of correct rotations, and the classifying threshold τ . The number of possible image combinations is $\binom{120}{n_A}$, making it computationally infeasible to test every combination, therefore images are randomly selected and tested 1 000 times. Because rotations are based on the randomly selected images, testing with a specified percentage of correct rotations is difficult as this percentage is based on the selected images. The correct rotation percentage (CRP) is calculated for each test as

$$\text{CRP} = \frac{\text{TP} + \text{FN}}{n_R}. \quad (25)$$

CRP can be increased by randomly selecting an initial pose, which forms a test set containing only the poses within a given angle of the initial pose. To ensure a high variability in the possible combinations, the initial pose is randomly selected 20 times, and if the test set contains fewer than 15 images, it is discarded and a new initial pose is selected. As the testing is performed over a total of 20 000 tests, the mean and standard deviation of CRP are given.

Matthew's correlation coefficient (MCC) [1],

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}}, \quad (26)$$

is used to provide a single measure of each testing scenario. It returns a measure between the verification and ground truth, where a value of +1 is perfect, 0 is no better than random and -1 is total disagreement. The four counts used by MCC are taken from all 20 000 tests, to produce a single value. It is not possible to produce an uncertainty for the MCC as individual tests may produce a division by zero; this occurs when all the verification or ground truth classifications are the same.

Testing is performed using $n_A = 4, \dots, 10$, with the MCC

value calculated for each threshold. The three angles used to adjust the CRP are given in Table 2. Using all 120 images is also tested as it produces a single fixed result.

Angle (deg)	Approx. CRP
30°	70%
70°	34%
No Restriction	17%

Table 2: The angles used to adjust the correct rotation percentage (CRP).

5. RESULTS

The main area of investigation is the accuracy of triplet verification. Efficiency is discussed in Section 6.

5.1 Accuracy

The accuracy of the verification classification is based on a variety of factors. To gauge how this accuracy changes, three factors are varied; the number of images, the percentage of correct rotations, and the classifying threshold. Table 3 lists the MCC values for each tested scenario. The factor that has the greatest influence on the verification classification is the CRP value. As the CRP value increases so does the accuracy of the verification classification. The classification ability of the three tested thresholds is closely linked to the number of used images. When using four images the thresholds perform almost identically, but as the number of images increases, only the MCC value of τ_b increases, with the MCC values of τ_a and τ_c typically decreasing. Of the three tested thresholds, τ_b provides the best classification. When all 120 images are used, the CRP cannot be increased, making the 30° and 70° inapplicable. The MCC value for τ_c is invalid when all images are used as TP + FP = 0.

The results for 6 and 10 images are shown with greater detail using receiver operating characteristic curves in Figure 5, where the classification threshold is varied between 0 and $n_A - 1$. These two graphs confirm the MCC values by showing that as the CRP increases, the area under the curve also has a significant increase. They also reveal how the thresholds behave, where as the number of images increases, τ_a and τ_c are pushed to the ends of the curves, with

No. of Images	Matthew’s Correlation Coefficient								
	30°			70°			All Images		
	τ_a	τ_b	τ_c	τ_a	τ_b	τ_c	τ_a	τ_b	τ_c
4	.44	.40	.40	.26	.22	.22	.16	.16	.16
5	.43	.49	.34	.29	.31	.21	.16	.19	.16
6	.39	.50	.32	.30	.30	.18	.16	.20	.14
7	.31	.50	.29	.31	.37	.16	.15	.23	.12
8	.30	.57	.28	.29	.33	.13	.15	.23	.11
9	.29	.60	.24	.26	.38	.12	.15	.25	.09
10	.23	.62	.25	.28	.37	.11	.14	.25	.08
120	N/A			N/A			.02	.33	–

Table 3: The Matthew’s correlation coefficient for all tested scenarios (rounded to 2 decimal places), where each scenario is determined by the number of images, the angle used to adjust the correct rotation percentage (30°, 70° and no restriction) and the classifying threshold (τ_a , τ_b or τ_c).

τ_b trending towards the centre of the curve, especially for large CRP.

6. DISCUSSION

Increasing the CRP is shown to also increase the classification ability of rotation verification, however, the method employed to vary the CRP in the performed tests is not directly applicable in practice. Improving the CRP of the SHCC is possible by either modifying the acquisition process to increase overlap, which is not always suitable, or by enhancing the cross-correlation algorithm to find a more ideal rotation. Enhancing the cross-correlation is not investigated as it is outside the scope of this paper.

Verifying that rotations conform is not limited to triplets of rotations, as the idea outlined in Section 3.1 for establishing consistency can be extended to four or more rotations. Depending on how the data have been acquired, using triplets for verification may not be appropriate, as there may not be a set of three images which overlap. Triplet verification was chosen as more triplets are formed from a set of images than any other group size, providing a greater level of insight into how each rotation should be classified as more cross-referencing can occur. Increasing the number of images will further bolster the amount of cross-referencing.

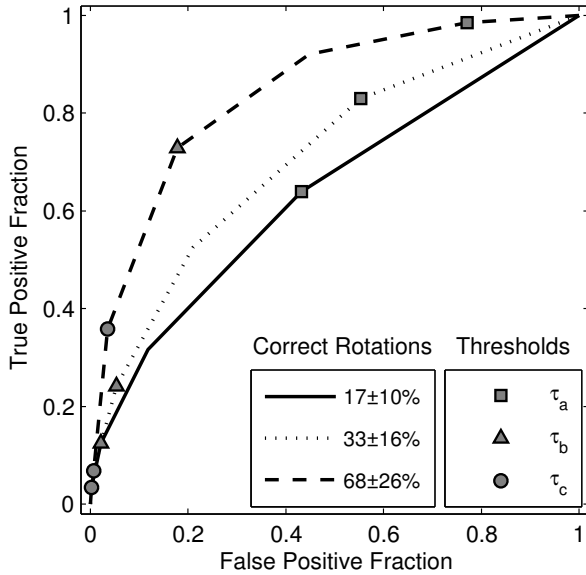
The choice of classification threshold has a significant impact on verification accuracy, especially when the image count increases. The threshold τ_a is excessively lenient, producing a high false-positive fraction; if triplet rotations were both consistent and correct, τ_a would work well, but as triplets can contain incorrect rotations whilst appearing consistent, it does not. Threshold τ_c is the most strict, producing a high false-negative fraction, and is only correct when all triplets that a rotation contributes to are consistent. Threshold τ_b avoids these two pitfalls as it only requires half of the triplets to be consistent; as the CRP and number of images increases, τ_b trends towards the optimal threshold location, as shown in Figure 5b. However, if the CRP is low, there is a greater number of incorrect rotations, decreasing the likelihood of triplet consistency. Incidentally, using four images causes both τ_b and τ_c to give the same threshold value, but differ when five or more images are used.

This paper focuses on the performance of triplet verification with respect to SHCC, therefore, the efficiency of SHCC is not investigated. The performance of triplet verification in contrast is essentially a non-issue due to the minimal amount of processing required. Classifying each triplet requires two 3-by-3 matrix multiplications, equation (13), followed by extracting the angular difference of the resulting matrix, equation (14). The expensive part of the verification is the $n_A(n_A - 1)(n_A - 2)/3$ matrix multiplications.


Applying triplet verification to a set of images does not provide a quintessential classification for a priori unknown rotations. Considering this, it is a beneficial addition to SHCC as the majority of its classifications are accurate. If triplet verification is used in conjunction with the translation correlation value, which indicates the quality of the alignment, the likelihood of selecting the correct rotations for aligning a set of images together is further improved.

7. CONCLUSION

Registering point-clouds using spherical harmonic cross-correlation does not guarantee correct rotational alignment.



(a) 6 Images



(b) 10 Images

Figure 5: Receiver operating curves showing how well rotation verification performs when the percentage of correct rotations increases. As the testing is performed on randomly chosen images, the percentage of correct rotations will fluctuate, therefore the standard deviation of this fluctuation has been provided.

Incorrect rotations occur when point-clouds have little overlap, which may cause the best found alignment from the correlation matrix to be erroneous. Using triplets of images, rotation verification is shown to be a viable method for classifying rotations, as it determines whether rotations are consistent with each other.

Classifying rotations using extreme thresholds is inadequate, as they are too strict or too lenient. A compromise threshold that requires a majority of correct triplets shows a trend towards optimal rotation classification, especially as the percentage of correct rotations and the number of registered images increases. Rotation verification is a beneficial addition to spherical harmonic cross-correlation, where even if the correct rotation percentage is low, individual rotations are still likely to be classified correctly.

8. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support provided by the University of Waikato Strategic Investment fund. Robert Larkins also acknowledges the financial support provided by both the Range Imaging and Waikato Doctoral Scholarships.

9. REFERENCES

- [1] P. Baldi, S. Brunak, Y. Chauvin, C. A. F. Andersen, and H. Nielsen. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*, 16(5):412–424, 2000.
- [2] F. Blais. Review of 20 years of range sensor development. *Journal of Electronic Imaging*, 13(1):231–243, January 2004.
- [3] G. E. Christensen and H. J. Johnson. Consistent image registration. *IEEE Transactions on Medical Imaging*, 20(7):568–582, July 2001.
- [4] X. Geng, D. Kumar, and G. E. Christensen. Transitive inverse-consistent manifold registration. *Lecture Notes in Computer Science: Information Processing in Medical Imaging*, 3565:468–479, July 2005.
- [5] B. Gutman, Y. Wang, T. Chan, P. M. Thompson, and A. W. Toga. Shape registration with spherical cross correlation. In *Second MICCAI 2008 Workshop on Mathematical Foundations in Computational Anatomy (MFCA '08)*, pages 56–67, 2008.
- [6] D. F. Huber and M. Hebert. Fully automatic registration of multiple 3D data sets. *Image and Vision Computing*, 21(7):637–650, July 2003.
- [7] R. L. Larkins, M. J. Cree, and A. A. Dorrington. Analysis of ICP variants for the registration of partially overlapping time-of-flight range images. In *Proceedings of the 25th International Image and Vision Computing New Zealand Conference (IVCNZ 2010)*, November 2010.
- [8] R. L. Larkins, M. J. Cree, and A. A. Dorrington. Analysis of binning of normals for spherical harmonic cross-correlation. In *Proceedings of SPIE-IS&T Electronic Imaging*, volume 8290, March 2012.
- [9] A. Pooja and V. M. Govindu. A multi-view extension of the ICP algorithm. In *Proceedings of the Seventh Indian Conference on Computer Vision, Graphics and Image Processing, ICVGIP '10*, pages 235–242, Chennai, India, December 2010.
- [10] J. Salvi, C. Matabosch, D. Fofi, and J. Forest. A review of recent range image registration methods with accuracy evaluation. *Image and Vision Computing*, 25(5):578–596, 2007.