

# Almost-Isometries for Point Cloud and Image Registration

Cyrus Anderson

July 2016

## 1 Introduction

Matching sets of points in 3D point clouds and aligning images are problems that are common to robotics, GIS, and medical imaging. The transformations resulting from the alignments can aid in areas ranging from robot localization to the measuring of what has changed between images, such as landscape developments or the healing after a surgery. The main contribution of this work is the implementation of an algorithm to find parameterized nonrigid transformations with the property being close to an isometry, that is, a rigid transformation. This algorithm is first formulated for labeled point clouds and then extended to the unlabeled case.

## 2 Related Works

The problem of finding a rigid transformation to align two sets of labeled points is well studied and known as the Procrustes problem. Having labeled point sets lacking an exact rigid transformation between them is a scenario that frequently arises from applying feature detectors to some underlying scene. The correspondences created by the detector may contain noise in the position of detected points, preventing a rigid transformation. In the case where noise in labeling is a larger problem, this scenario may be handled by RANSAC or other probabilistic methods [1, 2]. For handling the noise in points' position rather than labels, the proposed method may be used here to find a nearly rigid transformation. When we have unlabeled points and wish to find a nonrigid transformation, methods can largely be categorized as Iterative Closest Point (ICP) or energy based. Methods of the ICP family iteratively determine point correspondences and use those to determine a transformation [3, 4]. Energy based methods model the transformation of pixels or voxels as fluids and aim to minimize the energy of the system [5, 6, 7].

## 3 Framework

Let  $P = \{p_i\}_1^N$  and  $Q = \{q_i\}_1^N$  be point sets of corresponding points in  $\mathbb{R}^D$ . The goal of the Procrustes problem is to find  $\Phi : \mathbb{R}^D \rightarrow \mathbb{R}^D$  such that:

$$\forall i \quad \Phi(p_i) = q_i.$$

where  $\Phi$  is a rigid motion, that is,  $\Phi(x) = Rx + t$ ,  $R \in SO(D)$ , and  $t \in \mathbb{R}^D$ . For cases where there is noise in the points measured after a rigid transformation, or when the underlying transformation itself is nonrigid, we can instead aim to find a  $\Phi$  specifying a nonrigid transformation by minimizing the Euclidean distance between points. In order to prevent overfitting we can include a smoothing term controlled by an adjustable parameter  $\lambda$  that will enforce smoothness in the transform. We can then search for  $\Phi$  such that:

$$\Phi^* = \operatorname{argmin}_{\Phi} \sum_i (\Phi(p_i) - q_i)^2 + \frac{\lambda}{N^2} \sum_{i < j} \left( \frac{\|\Phi(p_i) - \Phi(p_j)\|}{\|p_i - p_j\|} - 1 \right)^2$$

Here  $\|\cdot\|$  denotes the  $L_2$  norm. The smoothing term checks that the transformation preserves pairwise distances, and the  $N$  is used to normalize for the number of points. The smoothing term also follows from the work in [8], to ensure that the transformation is almost a rigid motion. We now present different parameterizations of  $\Phi$ .

## Nonrigid Rotations in 2D

We first define the matrix that effects the nonrigid portion of the rotation:

$$S(x) = \begin{pmatrix} \cos(f(x)) & \sin(f(x)) \\ -\sin(f(x)) & \cos(f(x)) \end{pmatrix}$$

where  $f$  is chosen such that  $\forall t \in R \ \|tf'(t)\| < c\epsilon$  for some scalar  $c$  [8]. Throughout the rest of this paper, an admissible function  $f$  is chosen to be:

$$f(x; c) = c\sqrt{\log(1+x)}$$

The nonrigid rotation can then be written as:

$$\Phi_R(x; \mu, c) = R_\mu^T S(\|x\|)x$$

where  $R_\mu \in SO(2)$  and  $R_\mu$  is parameterized by the angle  $\mu$ .

## Nonrigid Translations in 2D

Translations can be parameterized with:

$$\Phi_T(x; c_1, c_2) = x + T(\|x\|), \quad T(x) = \begin{pmatrix} g_1(x) \\ g_2(x) \end{pmatrix}$$

where similar to  $f$  in  $\Phi_R$ , each  $g_i$  is chosen such that  $\forall t \in R \ \|g'_i(t)\| < c_i\epsilon$  for some scalar  $c_i$  [8].

## Compositions

In order to effect both a rotation and a translation, we can set:

$$\Phi_C(x; \mu, c_r, c_{t1}, c_{t2}) = \Phi_T \circ \Phi_R(x).$$

## Multiple Centers of Nonrigidity

Since the previous  $\Phi$  are all centered at the origin we can modify the parameterization to handle multiple centers that may not be at the origin. To shift the centers we can use:

$$\Phi_{SC}(x; \mu, c_r, c_{t1}, c_{t2}, c_1, c_2) = R_\mu^T S(\|x - \mathbf{c}\|)x + T(\|x - \mathbf{c}\|).$$

Note that the addition of  $\mathbf{c}$  adds two dimensions to the search space of the optimization routine. Now incorporating multiple centers, the equation becomes:

$$\Phi_{MC}(x) = \Phi_{SC_1} \circ \dots \circ \Phi_{SC_M}(x).$$

An example of a transformation with multiple centers is shown in figure 1. The points in  $\mathbb{R}^2$  are transformed to a greater degree near the centers of nonrigidity. Farther away from the centers, the nonrigid effect decays into to a rigid transform. This can be useful for applying nonrigidity to only a small area in the point set. When the function effecting the nonrigid rotation is parameterized with a small  $c$ , the amount of nonrigidity is smaller. In figure 2, the effect of a larger  $c$  is shown.

## 4 Optimization

Since the proposed  $\Phi$  are parameterized, we may write the formula as the minimization over those parameters, and solve the problem via a standard optimization routine (such as L-BFGS). It was found that the  $O(N^2)$  time needed to calculate the smoothing term was by far the most time consuming part of the the process, so the term was instead approximated. For every point, instead of calculating the pairwise distances to each of the other points, the distances to  $k$  other points were uniformly randomly picked. The normalization term is changed correspondingly, becoming  $kN$ . Using  $k = 4$  for each of the following experiments was found to be much faster.

## 5 Extension to the Unlabeled Case

We can extend the method to the unlabeled case by working with pixel or voxel values rather than their coordinates in the point set. Let an image  $I$  be defined as  $I : \mathbb{R}^D \rightarrow \mathbb{R}^K$ . For example, an RGB image would map from pixel coordinates in  $\mathbb{R}^2$  to values in  $\mathbb{R}^3$ . In this setting, a transformation  $\Phi$  modifies a source image  $I_p$  such that  $I_p \circ \Phi = I_q$ , with  $I_q$  being the destination image. Given source and destination images  $I_p$  and  $I_q$ , we can write the new optimization problem as:

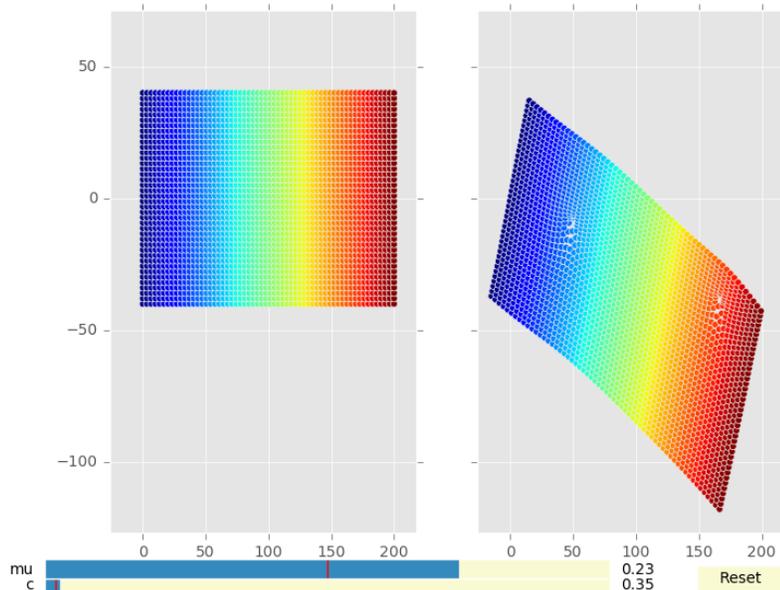


Figure 1: Nonrigid rotations with multiple centers. The original points in  $\mathbb{R}^2$  (*left*) are rotated, and closer to the centers pulled inwards (*right*). The parameters used to produce each center are:  $(0.23, 0.35, -15, 0, 0, 0)$ , and  $(0, 0.1, 0, 0, 170, 20)$ .

$$\Phi^* = \operatorname{argmin}_{\Phi} \sum_i (I_p(p_i) - I_q(\Phi^{-1}(p_i)))^2 + \frac{\lambda}{N^2} \sum_{i < j} \left( \frac{\|\Phi(p_i) - \Phi(p_j)\|}{\|p_i - p_j\|} - 1 \right)^2$$

Here, the pixel values between the source image and the destination image transformed to match the source are compared for all coordinates  $p_i$  of the source image. In practice,  $\Phi^{-1}$  may map a pixel coordinate to a coordinate between pixels in the mapped-to image. The value at that point can then be interpolated from nearby pixel values. We use bicubic interpolation in all of the evaluations.

This new formulation can be susceptible to local minima. In order to improve convergence, a coarse to fine regime is used. Rather than immediately finding  $\Phi^*$  for the original image pair,  $\Phi^*$  is first found for blurred versions of the original images. This  $\Phi^*$  is used as the starting point for the next optimization, which finds  $\Phi^*$  for a less blurry pair. This iteration continues until  $\Phi^*$  for the original image pair is found. Gaussian blur with successive kernel sigmas of 5, 4, 3, and 1.5 are used in the evaluations.

## 6 Evaluation

For both the labeled and the unlabeled problems, we compare the proposed methods to rigid motions in the sum of squared errors (SSE) over the pixel values of the rectified transformed image

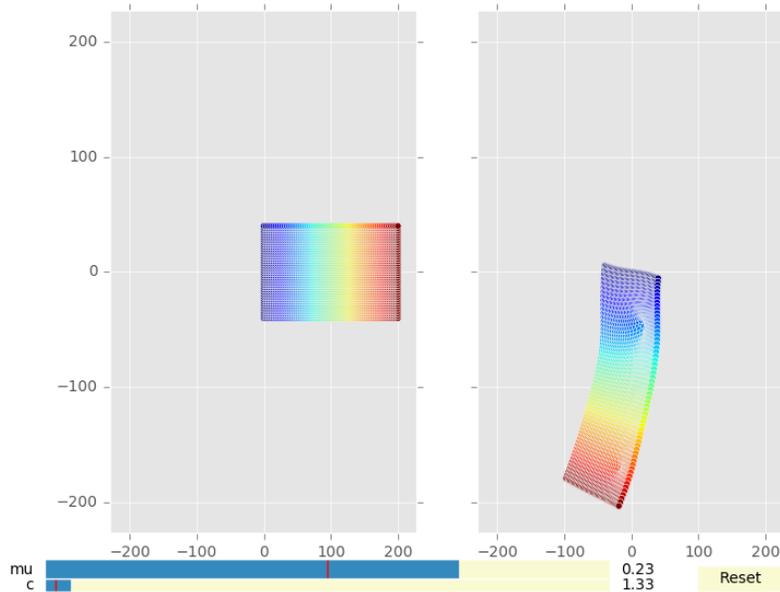


Figure 2: Nonrigid rotations with multiple centers. The original points in  $\mathbb{R}^2$  (*left*) are transformed with two centers of nonrigidity (*right*). The blue center’s greater nonrigidity than the red center pulls the points from farther away, creating a sort of bend in the sheet of points. The parameters used to produce each center are:  $(0.23, 1.33, -15, 0, 0, 0)$ , and  $(0, 0.1, 0, 0, 170, 20)$ .

and the original image pairs. The nonrigid transformations used for the evaluations consist of two centers and smoothed with a  $\lambda$  of 1.

## 6.1 Point Cloud Based Registration

Each pixel coordinate in the image lending itself to a point, we first construct the transformed image with a known transformation on the coordinates to obtain the point cloud after the transformation. The transformation learned between the two point clouds is then used to warp the images, allowing for the calculation of the SSE. The results are shown in table 1. The rigid and nonrigid methods produce fits with some difference in SSE, but looking at figures 4 and 5, we see that they perform very similarly.

## 6.2 Pixel Based Unlabeled Registration

The unlabeled formulation works directly on the images here. In order to improve convergence, in addition to the coarse to fine regime, the rigid transformation is applied before finding the nonrigid transformation. The nonrigid registration method outperforms the rigid registration method, and both outperform the methods applied to the labeled problem (*table 1*). Comparing figures 6 and 7, we see that the nonrigid transformation overcorrects for the deformity in the triangle in order

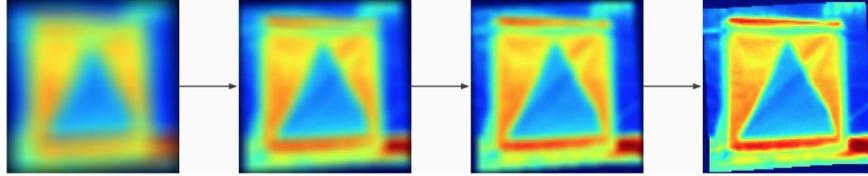


Figure 3: Avoiding local minima with blur. From right to left, the Gaussian blur applied to the original image is decreased until we are left with the original.

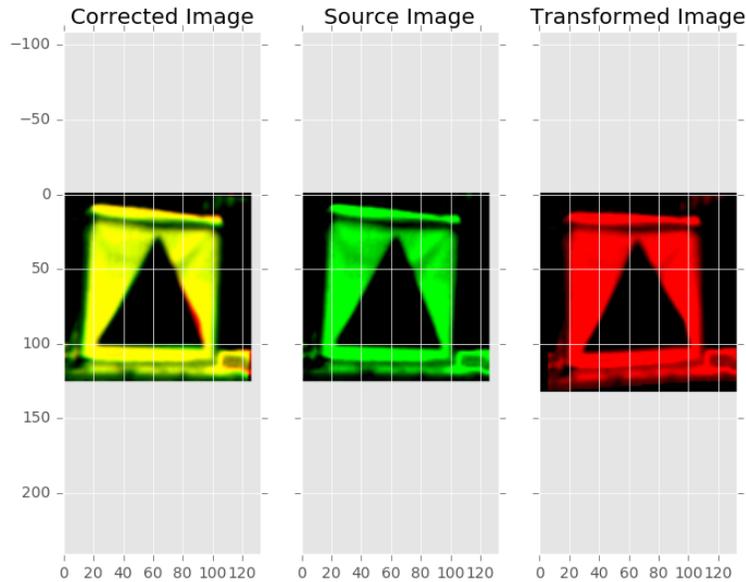


Figure 4: Rigid registration with point cloud information. The original image (*center*) is transformed to the image on the right, and the corrected version of the transformed image is overlaid on the original (*left*).

to match more pixels at the boundary of the image. This gain in boundary pixels offsets the larger deformity in the triangle, explaining the decrease in SSE.

## 7 Conclusion

We presented a method for registering labeled point clouds with a nonrigid deformation, and extended the formulation to the unlabeled setting for images. The nonrigid model is not fit as easily as rigid transformations, but its additional flexibility allows it to achieve lower registration error. Future avenues of research could include the use of specialized optimization techniques to improve the fitting process, and an automated method of determining the optimal number of centers used in the nonrigid transformations.

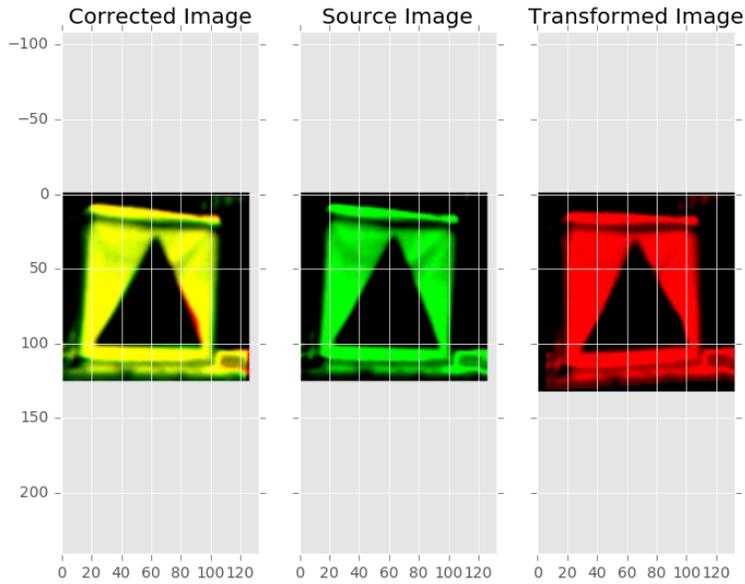


Figure 5: Nonrigid registration with point cloud information. The original image (*center*) is transformed to the image on the right, and the corrected version of the transformed image is overlaid on the original (*left*).

## Acknowledgments

The author thanks David Speyer, Michael Lu, Neo Charalambides, Sean Kelly, and Steven Damelin for their fruitful comments and discussions. Support from the NSF is gratefully appreciated.

## References

- [1] M. A. Fischler and R. C. Bolles, “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography,” *Communications of the ACM*,

Registration Method	SSE
Labeled, Rigid	382.648
Labeled, Nonrigid	382.857
Unlabeled, Rigid	362.129
Unlabeled, Rigid with Nonrigid	361.519

Table 1: SSE for the different methods. The methods solving the unlabeled registration problem outperform the labeled methods, and the extra freedom in the nonrigid model leads it to outperform the rigid registration in the unlabeled setting.

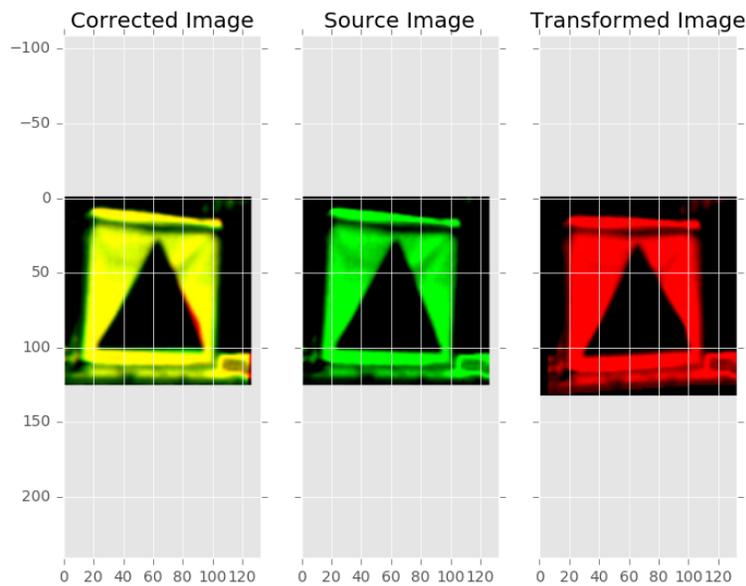


Figure 6: Rigid registration with pixel information. The original image (*center*) is transformed to the image on the right, and the corrected version of the transformed image is overlaid on the original (*left*).

vol. 24, no. 6, pp. 381–395, 1981.

- [2] V. M. Govindu and M. Werman, “On using priors in affine matching,” *Image and Vision Computing*, vol. 22, no. 14, pp. 1157–1164, 2004.
- [3] J. Ma, J. Zhao, J. Tian, Z. Tu, and A. L. Yuille, “Robust estimation of nonrigid transformation for point set registration,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2147–2154, 2013.
- [4] D. Haehnel, S. Thrun, and W. Burgard, “An extension of the icp algorithm for modeling nonrigid objects with mobile robots,” in *IJCAI*, vol. 3, pp. 915–920, 2003.
- [5] J. Rexilius, S. K. Warfield, C. R. Guttmann, X. Wei, R. Benson, L. Wolfson, M. Shenton, H. Handels, and R. Kikinis, “A novel nonrigid registration algorithm and applications,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 923–931, Springer, 2001.
- [6] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, “Computing large deformation metric mappings via geodesic flows of diffeomorphisms,” *International journal of computer vision*, vol. 61, no. 2, pp. 139–157, 2005.
- [7] L. Risser, F.-X. Vialard, H. Y. Baluwala, and J. A. Schnabel, “Piecewise-diffeomorphic image registration: Application to the motion estimation between 3d ct lung images with sliding conditions,” *Medical image analysis*, vol. 17, no. 2, pp. 182–193, 2013.

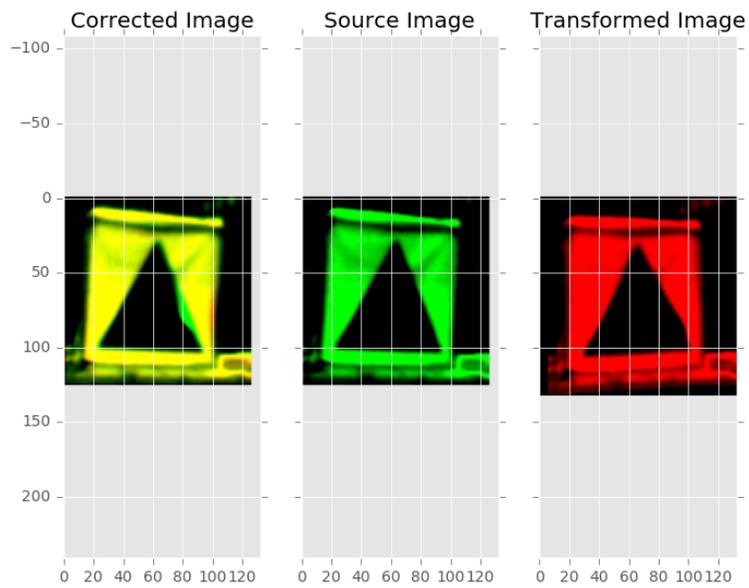


Figure 7: Nonrigid registration with pixel information. The original image (*center*) is transformed to the image on the right, and the corrected version of the transformed image is overlaid on the original (*left*).

- [8] S. Damelin and C. Fefferman, “Extensions, interpolation and matching in  $\mathbb{R}^D$ ,” *arXiv preprint arXiv:1411.2451*, 2014.