
The Use of Robust Local Hausdorff Distances in Accuracy Assessment for Image Alignment of Brain MRI

Release 0.10

Eric Billet, Andriy Fedorov and Nikos Chrisochoides

May 2, 2008

Center for Real Time Computing
Department of Computer Science
College of William and Mary
Williamsburg, VA 23185
{ejbill,fedorov,nikos}@cs.wm.edu

Abstract

We present and implement an error estimation protocol in the Insight Toolkit ([ITK](#)) for assessing the accuracy of image alignment. We base this error estimation on a robust version of the Hausdorff Distance (HD) metric applied to the recovered edges of the images. The robust modifications we introduce to the HD metric significantly reduce the amount of outliers in the local distance error estimation. We evaluate the accuracy of our protocol on synthetically deformed images. We provide the source code and datasets to reproduce this evaluation. The proposed method is shown to improve error assessment when it is compared with conventional HD methods. This approach has many applications including local estimation and visual assessment of registration error and registration parameter selection.

Contents

1	Introduction	2
2	Methods	3
3	Experimental Framework	4
4	Results	6
4.1	3D Image Results	6
4.2	2D Image Results	8
5	Implementation	9
5.1	itk::LocalDistanceMapImageFilter	10
5.2	itk::LocalGrayscaleDistanceMapImageFilter	10
5.3	itk::LocalDistanceMapSmoothingImageFilter	12
5.4	RunAssessment.cxx	12

5.5	DFGenerator.cxx	13
5.6	Evaluation.cxx	13
5.7	Software Requirements	13
6	Conclusions	13
7	Acknowledgments	14

1 Introduction

Assessing the accuracy of image alignment is a vital step in image registration. Accuracy assessment is essential for verification of registration quality, comparison of registration results, and registration parameter selection. However, estimating registration accuracy is a challenging problem for the medical imaging community. Clinical application of registration methods require the validation of registration results, but this validation can only be performed on cases where ground truth is known (cadavers, invasive markers, phantoms). The ability to assess registration quality of non-rigid brain registration during image guided neurosurgery is limited; an expert can assess quality by locating a series of clearly identifiable landmarks or regions in both intra-operative and pre-operative images. Arguably, this may provide the best precision, but the assessment is limited to the easily identifiable landmarks. This method is also time consuming and requires the involvement of an expert which may not be practical during neurosurgery, when comparison of thousands of registrations is necessary. Consider a scenario when the optimum parameters of a registration method are not known, and we need to perform a large scale parametric search for the best parameter configuration. In this case, we are facing comparison of hundreds and even thousands of registration results, which in practice is only feasible using automatic or semi-automated approaches. We attempt to automate this assessment process through the implementation of local similarity metrics.

Although, similarity metrics are no substitute for validation, similarity metrics have been used to assess registrations. They can provide confidence in the registration result as well as criteria for comparison between registration methods and parameters. *One common shortcoming of this approach is that similarity metric values do not provide quantification in terms of distance in physical space.* Thus, one has the ability to track improvements among registrations but cannot speculate about the degree of misalignment

One approach to this problem is the use of Hausdorff Distance (HD) metric. The HD metric is derived from Euclidean distances between point sets and thus related to local alignment distances unlike many other image similarity measures. The original HD metric is highly sensitive to noise [6], but variations of the metric have been proposed to improve the robustness of the metric. A survey of these modifications can be found in [13, 3]. Even with the correspondence to physical space, HD based metrics have scarcely been used for assessment of image alignment. They were used by Morain-Nicolier et. al [8] to quantify brain tumor evolution and later by Archip et al. [1, 2] to assess performance of non-rigid image registration of brain MRI. To our knowledge no formal study has been completed to investigate the use of HD as an accuracy assessment measure for image alignment. We evaluate existing metrics and implement an automated assessment approach based on a robust variation of the HD metric. We provide three ITK filters for local assessment and implement an evaluation framework for comparison of these methods. Based on our findings, our robust HD approach can significantly improve the results as compared to the previously used version of HD.

2 Methods

Problem: assess accuracy of alignment between two images, I and J .

We develop a method to solve this problem through estimation of point-wise alignment error. The HD metric operates on sets of points, and in order to apply it to the problem we need first to identify such (feature) points in the images. Consistency of feature points identification is essential to the robustness and correctness of the metric. Thus, prior to local estimation, several preprocessing steps must be employed on the input images.

First, we smooth the input images using anisotropic diffusion [9] to filter out noise. Anisotropic diffusion smoothing is chosen because it preserves features and all of the similarity metrics we examine are based on comparison of features of images rather than intensities. Second, adaptive contrast equalization [12] is performed to improve the edge detection results. Third, edges are identified with a Canny edge detection filter [4]. Thresholds for Canny edge detection are adaptively selected using binary search to have approximately the same number of edges in each image. The number of edge points is controlled by changing the upper and lower values of the threshold parameters in the edge detector [4]. All preprocessing steps were implemented using ITK filters. The implementation details can be found in Section 5.

Once the feature edges (points) are identified in both images, we can depart from the images and analyze the alignment accuracy by working with the two sets of points. Let A and B be the set of binary images created from extracting features from I and J , respectively, and $A = \{a_1, a_2, \dots\}$ and $B = \{b_1, b_2, \dots\}$ be the corresponding set of feature points (non-zero pixels). The directed HD between these two sets of points, $h(A, B)$, is defined as the maximum distance from any point in A to any point in B . The symmetric HD, $H(A, B)$ is the maximum of both direct distances [6]:

$$h(A, B) = \max_{a \in A} (d(a, B)), \text{ where } d(a, B) = \min_{b \in B} \|a - b\| \quad (1)$$

$$H(A, B) = \max(h(A, B), h(B, A)) \quad (2)$$

HD provides a global comparison of similarity, and was used by Archip et al. [1, 2] to estimate the alignment of brain MRI images after non-rigid registration. However, the notion of local similarity is lost. For example, HD values can remain constant whether a set of images are misaligned in just one region or many regions. The local distance map (LDM) proposed by Baudrier et al. [3] is more suitable for point-wise error estimation.

The LDM derives a local measure of dissimilarity for 2D binary images by comparing images locally with a sliding window, where $H(A, B)$ is used as a dissimilarity measure within the window.




Figure 1: Binary image (gray/white cells) with corresponding grayscale image (D=3) after non-rigid registration.

$$\forall x \in \mathbf{R}^3 : H_{loc}(x) = |\mathbf{1}_{\mathcal{A}(x)} - \mathbf{1}_{\mathcal{B}(x)}| \times \max(d(x, A), d(x, B)), \quad (3)$$

$$\mathbf{1}_{\mathcal{A}x} = \begin{cases} 1 & \text{if } \mathcal{A}(x) \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{where } \mathcal{A}(x) \text{ is the voxel value of } A \text{ at location } x.$$

Remark 1. The conventional $H(A, B)$ can be derived by taking the maximum of the local, H_{loc} , comparisons. i.e., $H(A, B) = \max(H_{loc}(A, B))$

LDMaP is parameterless and yields a more descriptive measure of *local* similarity than HD, but the LDMaP as described by Baudrier et al. [3] does not improve robustness of local estimation. Ideally, we are looking for a measure which is equal or analogous to the distance between corresponding points in the two images. However, this cannot be assumed for HD, since HD has no point-to-point correspondence.

In this paper we extend the definition of LDMaP by adding the notion of correspondence by using a grayscale modification of the HD. Grayscale HD was proposed by Zhao et al. [13] for matching 2D images corrupted by noise. We extend this grayscale HD definition to 3D for medical image comparison and include it within the LDMaP definition.

Given two input binary images, A and B , let \tilde{A} and \tilde{B} be the grayscale images whose voxel values are initialized to the number of non-zero voxels in constant size neighborhood of the corresponding binary image. \tilde{A} and \tilde{B} have the same size and boundary as the initial binary images. The neighborhood is isotropic and of size $D \times D \times D$. Figure 1 provides a 2D example with $D = 3$. This modification improves the sense of point correspondence while calculating point distances because the corresponding point in the second image is likely to have similar or same number of neighbor feature points. The directed distance $d(a_g, \tilde{B})$ where g is the grayscale value at voxel a is defined as the minimum distance from point a to any point in \tilde{B} with grayscale value within some tolerance, t , of g . The tolerance allows to adjust the sensitivity to differences in edge images (i.e., from resampling, different imaging device, etc.).

$$d(a_g, \tilde{B}) = \min_{b_{g'} \in \tilde{B}} \|a_g - b_{g'}\|, \quad g - t \leq g' \leq g + t \quad (4)$$

Using Equation 4 we define the local grayscale HD, GH_{loc} as follows:

$$\forall x \in \mathbf{R}^3 : GH_{loc}(x) = |\mathbf{1}_{\mathcal{A}(x)} - \mathbf{1}_{\mathcal{B}(x)}| \times \max(d(x_g, \tilde{A}), d(x_g, \tilde{B})). \quad (5)$$

Similar to the Hausdorff Distance, Grayscale Hausdorff Distance (GHD) $GH(A, B)$ is defined as the maximum of the local calculations, $GH(A, B) = \max(GH_{loc}(A, B))$.

In addition, we further improve the robustness of Grayscale Hausdorff Distance by applying robust statistics. Let $RGH_{loc}(x)$ be the local Robust Grayscale Hausdorff Distance of images A and B at voxel x , and let W_x be an isotropic window of size $S \times S \times S$ centered around x . $RGH_{loc}(x)$ is defined as the robust average calculated from the voxels inside W_x . Least trimmed squares [11] is a reasonable choice for this robust average, but other averages can be employed to improve robustness. We define the Robust Grayscale Hausdorff Distance (RGHD) $RGH(A, B) = \max(RGH_{loc}(A, B))$.

The complete diagram of the alignment accuracy assessment framework, together with the parameters used at each step, is depicted in Figure 2.

3 Experimental Framework

We evaluate the effectiveness of the proposed accuracy assessment methodology for non-rigid registration of brain MRI using the synthetic ground truth data. First we use the method described by Rogelj et al. [10] to construct a synthetic deformation field. The synthetic deformation field is applied to the original grayscale image, followed by feature detection step performed on both the original and deformed images. The proposed error recovery methodology is then employed to estimate the misalignment between the original and

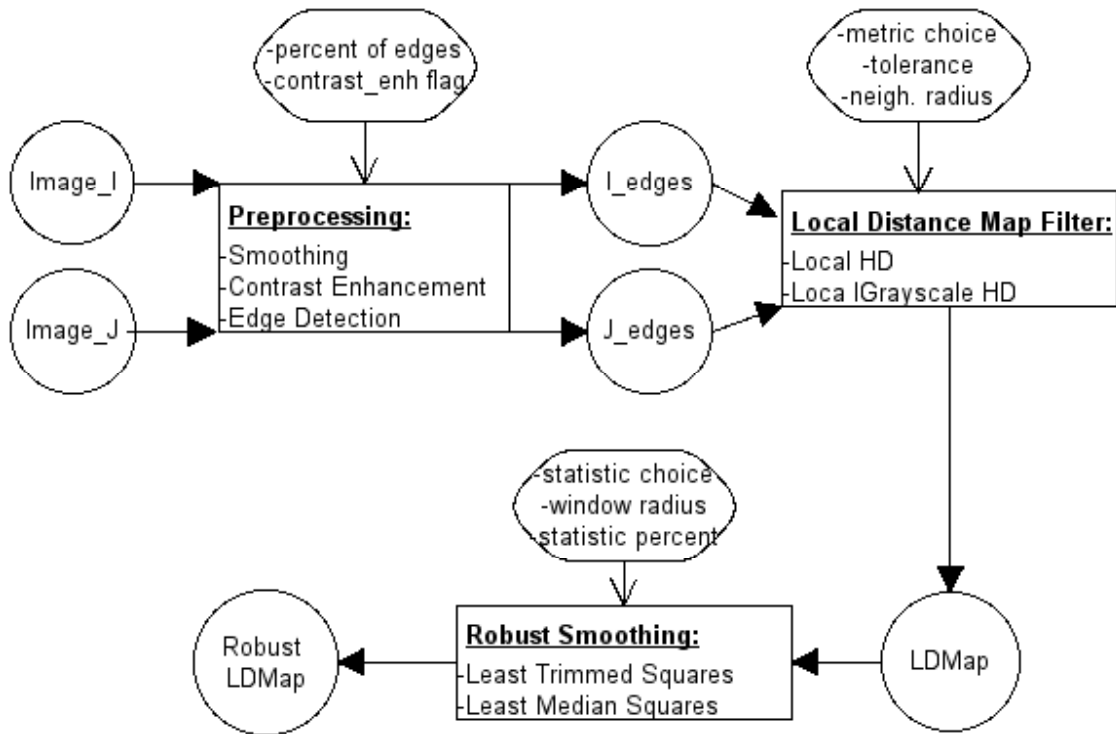


Figure 2: The basic components of the assessment framework: (1) preprocessing of input images, (2) computation of local distance metric on edge images and (3) robust smoothing of local distance map.

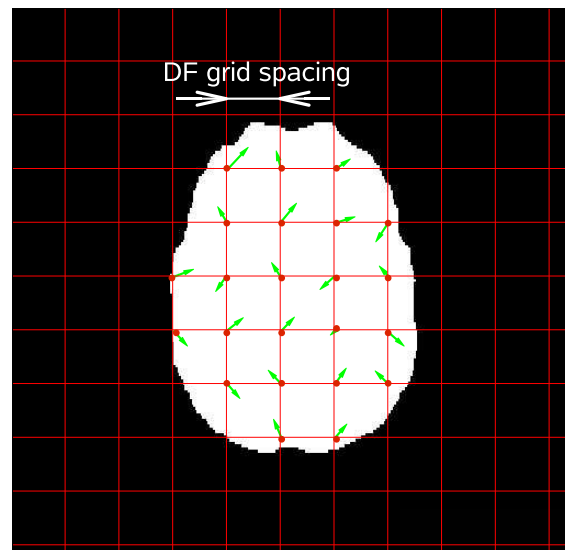


Figure 3: Random deformation vectors are generated at the knots (red circles) of the deformation field grid that are located within the user-defined mask.

the deformed images at the selected feature points. The true misalignment value, which is the magnitude of the synthetic deformation at a point, ideally should be equal to the error value recovered by the either of the proposed assessment measures.

In order to create the synthetic deformation field, we first construct a sparse point sample at the knots of an isotropic sampling grid. This procedure is illustrated in Figure 3, where point sample is shown with red circles. We assign a random deformation vector at each of these points, with the components of the vector drawn from a Gaussian distribution parametrized by mean($\mu = 0$) and variance(σ). The dense deformation field is constructed by using Thin Plate Splines interpolation at non-knot image points within the binary mask.

The current implementation of the synthetic deformation is not provided as a generic class. Instead, we have separate tools to construct 2D and 3D deformation fields. The deformation fields are parametrized by the sampling grid spacing and the variance of the Gaussian distribution, which should be provided by the user, together with the binary mask of the region of interest. `itk::ThinPlateSplineKernelTransform` is used to interpolate the values of the dense deformation field at non-grid pixels

We compare the local estimation methods with ground truth error using two measures: distribution of error and percentage of outliers. Ideally, the distribution of local error estimates (H_{loc} , GH_{loc} and RGH_{loc}) will closely mimic the true error distribution. Also, a good estimate of error should have minimum number of outliers. Let d_i be the local distance at point location i measured by a local estimation method and e_i be the true error at the same point. We define outliers as any point i where $|d_i - e_i| > 2mm$. We choose 2 mm because the deformation field is in physical space and the HD distance implementation is limited to 1 mm image spacing. Thus, errors as large as $\sqrt{3}$ cannot be prevented.

4 Results

We present results for each of the described metrics (HD, GHD, RGHD) using the evaluation framework described in Section 3 for both 3D and 2D images. The 3D images are not provided with this submission due to possible copyright issues, but can be freely downloaded from BrainWeb Simulated Brain Database [5, 7]. We used the following settings for constructing the simulated image: Modality=T1, Slice thickness = 1mm, Noise = {0%,9%} and RF=0%. The 2d images accompanying this paper were taken from the example data provided with ITK (ITK/Examples/Data/BrainProtonDensitySliceBorder20.png, BrainProtonDensitySliceBorder20Mask.png).

4.1 3D Image Results

The 3D results presented in this paper are produced using the following parameters. The variance parameter, σ , of the Synthetic Gaussian deformation was set to the integer values between 1 and 12 to get increasingly more complex deformations. The deformation grid spacing was set at 30 mm for all test cases. These settings create displacements with a maximum of 12 mm and an average between 1 and 4 mm, depending on variance. Parameters for GHD include: neighborhood size, for creation of grayscale images, set to $D = 3$ and tolerance for $GHD_{loc}(x)$, $t = 2$. RGHD used a least trimmed squares robust metric with percentage set to 80% and window size, $W_x = 11$. We are able to incorporate the use of noisy images in our evaluation because Brainweb provides a simulation of noisy images, and thus we evaluate local alignment assessment between undeformed images with 0% noise and deformed images with 0% and 9% noise.

The error distributions for: the actual error, H_{loc} , GH_{loc} , and RGH_{loc} for $\sigma = 5$ of the synthetic deformation

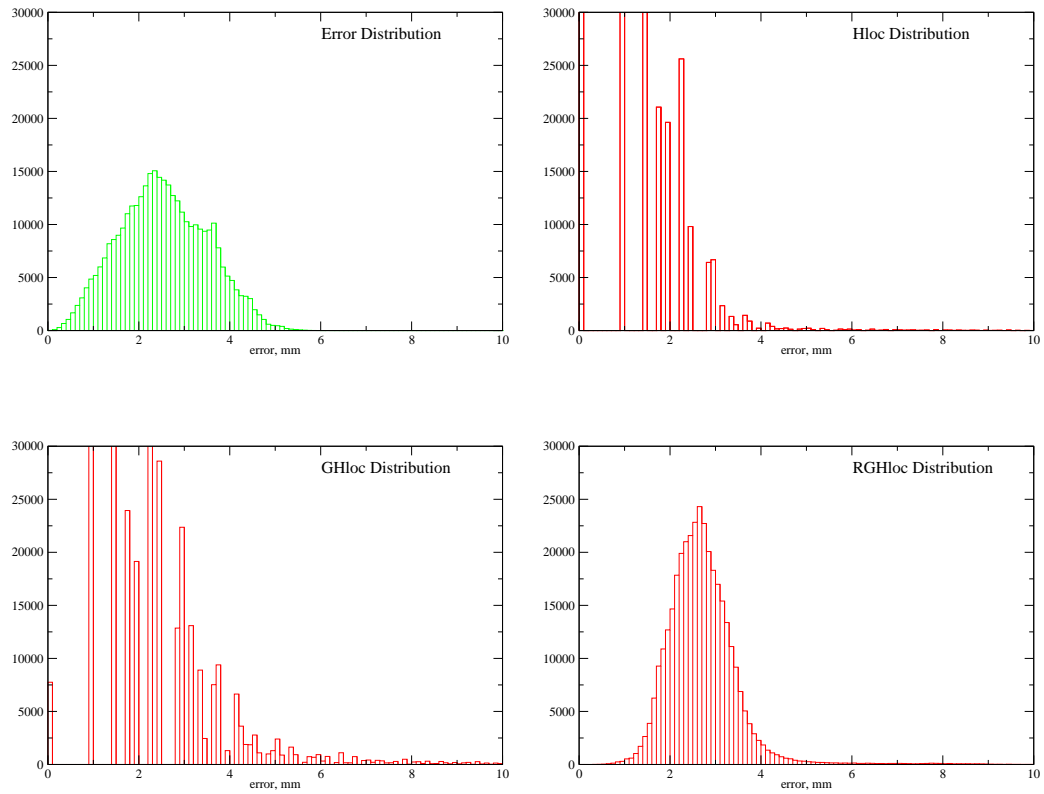


Figure 4: Distribution of the error and the HD, GH, and RGH values for the same synthetic deformation case (3D BrainWeb image, Gaussian deformation, variance 5).




Figure 5: The change in outlier percentages depending on the variance of Gaussian Deformation. Left: deformed 2D image included with paper Right: deformed 3D Brainweb images with 0% and 9% noise.




Figure 6: Global error statistics (i.e. HD, 95%HD, RGHD) for synthetically deformed 3D BrainWeb images with and without noise. RGHD results were derived from RGH_{loc} values.

are shown in Figure 4. Other variances had similar results. The distribution of RGH_{loc} is an improved approximation of the error distribution over H_{loc} and GH_{loc} .

As expected the data in Figure 5 indicates that RGH_{loc} is also the most robust method. Figure 5 displays the percentage of outliers in the set of local distance estimations. RGH_{loc} remains at a low percentage of outliers throughout the 3D testing. GH_{loc} also provides improved robustness over the H_{loc} estimations. For best results, we recommend the use of RGH_{loc} when deformations of magnitude 3mm to 7mm are present.

The error distribution and outliers help to build confidence in the robustness and accuracy of the local estimates. Next we examine the use of robust local estimation methods for non-rigid registration of brain MRI comparison. Our hypothesis is that summary statistics of local estimations provide a more robust global statistic for comparison purposes. We test this hypothesis by examining the 95% percentile of the local estimates at each variance and compare it with the results obtained with the conventional HD and the 95% PHD used by Archip et al. [1, 2]. These results are presented in Figure 6. RGHD tracks the improvement well (i.e. increasing constantly w.r.t. error) even in the presence of noise. Moreover, RGHD values consistently over-estimate (upper-bound) the true mean error. This is expected because the robust metrics are based on the assumption that the majority of outliers are underestimates and thus the calculated mean increases. The use of other summary statistics (RMS, LTS) on the local estimates may improve this discrepancy, but the study of these statistics is left to future work.

4.2 2D Image Results

We evaluate alignment assessment between the undeformed and deformed 2D ITK images with the following parameters. The deformation field is created with: $\sigma = \{1, 2, 3, \dots, 12\}$ and spacing = 30mm. Deformation vectors with magnitude averaging between 1.5mm to 7mm are observed, magnitudes increase with variance. A maximum deformation of 9mm occurs at $\sigma = 12$. The parameters for local estimation parameters are based on the 3D results: neighborhood of GH_{loc} , $D = 3$, tolerance of grayscale, $t = 1$, window size of RGH_{loc} , $W_x = 11$, and percentage = 80%. Adjustments of these parameters may cause different results. A more exhaustive study of parameter selection will be performed in the future and reported in upcoming updates of this document.

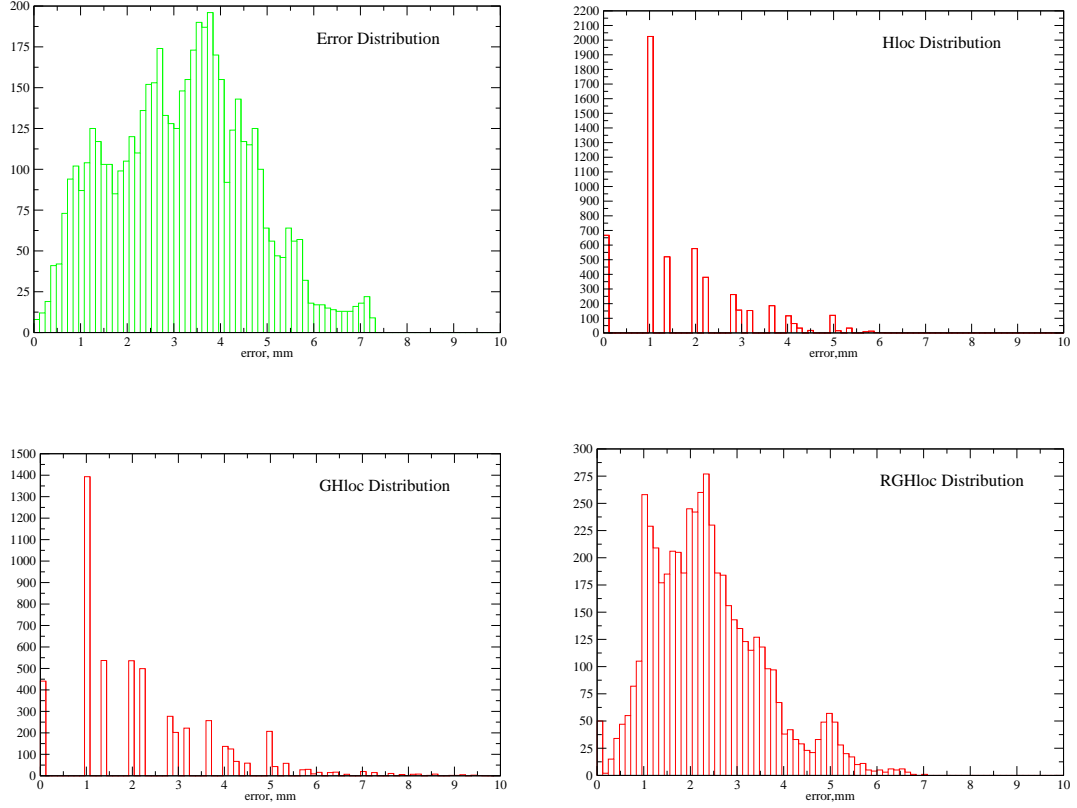


Figure 7: Distribution of the error and the HD, GH, and RGHD values for the same synthetic deformation case (2D ITK image, Gaussian deformation, variance 8).

Similar to the results in 3D data, RGHD provides the closest approximation to true error distribution with 2D images, see Figure 7. However, as seen in Figure 5 the outliers for RGHD are not as low with the 2D images. Many factors could contribute to this difference including the parameters chosen and the deformation field itself. Although, the variances were set to the same values, the deformation magnitudes are slightly higher on average in 2D image tests. The RGHD is still shown to be the most robust method compared.

We also examined the use of local estimation for visualizing the error in Figure 8. We color the LDMap created by RGH_{loc} estimations, allowing visualization of the magnitude of error throughout the image.

5 Implementation

The presented local error estimation methods have been implemented in three n-dimensional image to image filters: `itk::LocalDistanceMapImageFilter`, `itk::LocalGrayscaleDistanceMapImageFilter`, and `itk::LocalDistanceMapSmoothingImageFilter`. In addition, the entire alignment assessment process is implemented in `RunAssessment.cxx`. In Figure 5 we present the entire framework for our evaluation of local assessment methods. Descriptions of each of these filters and their use is described in the following sections. Two helper classes were also implemented. `itk::CountImageFilter` assists in the conversion of a binary image to its grayscale counterpart by implementing a simple counter of non-zero voxels in a predefined neighbor-




Figure 8: Local estimation of misalignment using RGHD. Left: Undeformed 2D image. Center: Deformed image, Gaussian variance of 8 mm. Right: LDM of the deformed and undeformed images using RGHD, color-coded, mm.

hood region. `itk::errorStatistics` is used to calculate descriptive error statistics from the robust LDM image, see Figure 2.

5.1 itk::LocalDistanceMapImageFilter

`itkLocalDistanceMapImageFilter` is the ITK implementation of local HD metric as described in Section 2. This filter is derived from the `itk::ImageToImage` filter and is templated over the two input images and the output image. This filter requires two distance maps as input. It outputs a local distance map with values of each voxel, x , equal to $H_{loc}(x)$.

This processing step does not have any parameters, and the only methods of relevance are those that provide the set and get functionality for the filter inputs:

- `SetInput1(const TDistance1 *)`, `GetInput1()`
- `SetInput2(const TDistance1 *)`, `GetInput2()`

The class is multi-threaded and the entire algorithm is contained in `ThreadedGenerateData()`.

5.2 itk::LocalGrayscaleDistanceMapImageFilter

We implemented a local grayscale Hausdorff filter based on the definition in Section 2. Similar to the previous filter, `itk::LocalGrayscaleDistanceMapImageFilter` is derived from the `itk::ImageToImage` filter and is templated over its two input images and one output image. This filter takes two binary edge images as input and outputs a local distance map with values of each voxel, x , equal to $GH_{loc}(x)$.

The first step is the construction of the grayscale image from the input binary image, as shown in Figure 1. This step is facilitated by the helper class `itk::CountImageFilter` we implemented, and is parametrized by the neighborhood radius value. As described in Section 2, to find the $GH_{loc}(x)$ values we must perform a search at each feature point, for the nearest neighbor features with values within a given range, $g - t \leq g' \leq g + t$. Exhaustive search is computationally expensive, instead the $GH_{loc}(x)$ values are computed by performing an

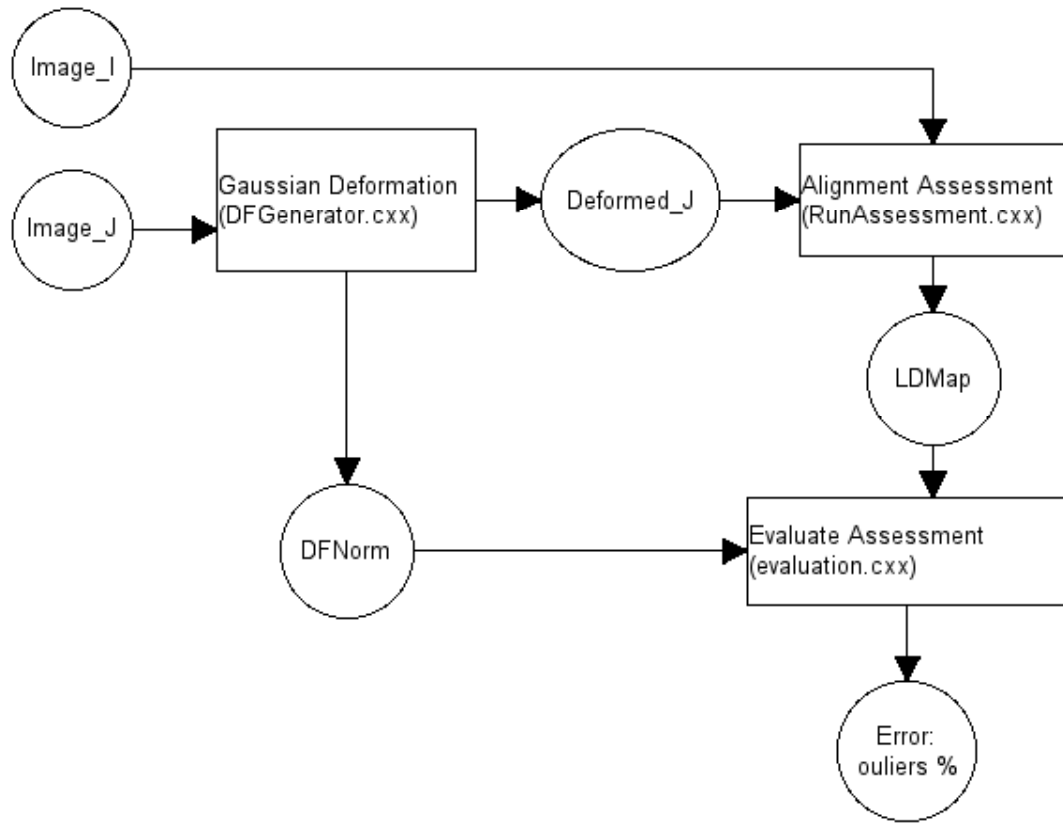


Figure 9: Components for evaluation of accuracy alignment.

expanding search with maximum radius, `MaxDef`, which is another parameter of the filter. `MaxDef` should be chosen based on the estimated maximum deformation value expected. The indices for the expanding search are determined in `BeforeThreadedGenerateData()`.

The input methods for this filter are:

- `SetInput1(const TEdge1 *)`, `GetInput2()`
- `SetInput2(const TEdge1 *)`, `GetInput2()`

The following parameters for this filter can be manipulated through Set/Get functions:

- **MaxDef**: This controls the search radius for finding the nearest neighbor, should be set to the maximum deformation possible (e.g. 10% of adult brain size). If no feature point is found within the search radius, the voxel is set to an unrealistic value (-100).
- **Tol**: the tolerance, t , for the local grayscale Hausdorff metric. Higher values of t , will create smaller distances, but the distinctiveness of corresponding points will decrease.
- **Radius**: radius of the neighborhood size, D , ($2 \times \text{radius} + 1 = D$), for which the grayscale images are created. Typical values used are 1 to 4.

The class is multi-threaded. The generation of grayscale images is performed in `BeforeThreadedGenerateData()`, but the rest of the calculations for local Grayscale HD is contained in `ThreadedGenerateData()`.

5.3 itk::LocalDistanceMapSmoothingImageFilter

itk::LocalDistanceMapSmoothingImageFilter is a smoothing filter designed specifically for post-processing of the local distance map to obtain the values of RGH_{loc} as described in Section 2. This filter applies a robust average within a window centered around a feature point. The filter only takes into account feature points. It accomplishes this by using “edge” (feature point) images as a mask during computation. The local robust grayscale Hausdorff distances (RGH_{loc}) are computed by applying this filter to the local distance map created by itk::LocalGrayscaleDistanceMapImageFilter.

The input methods for this filter are:

- SetInput1(const TEdge1 *), GetInput1() the first edge image, typically used to create local distance map
- SetInput2(const TEdge2 *), GetInput2() the second edge image, typically used to create local distance map
- SetInput3(const TDistance1 *), GetInput3() the local distance map

The parameters for this filter are manipulated by Get/Set commands:

- **Radius:** the radius, $(2 \times radius + 1 = S)$, of the window used for robust averaging
- **MinElements:** the minimum number of feature points needed to compute average. If below minimum assume there is no confidence in region and voxel set to unrealistic value (-100).
- **Percent:** the percentage used for the robust statistic, (i.e. $|1 - percent|$ is discarded in least trimmed squares)

Function GetRobustStat(std::vector) computes the robust statistic on the distances inside the window. Currently, only LTS is implemented, but other robust statistics could easily be added. As with the other filters, this filter is multi-threaded and derived from the ImageToImage filter.

5.4 RunAssessment.cxx

An example code for running the complete alignment assessment process on two images is provided in RunAssessment.cxx. Three basic components of the process (preprocessing, local distance map generation, robust smoothing) are each implemented. The preprocessing is performed using three ITK filters: itk::CannyEdgeDetectionImageFilter, itk::CurvatureAnisotropicDiffusionImageFilter, and itk::AdaptiveHistogramEqualizationImageFilter. Smoothing and contrast enhancement parameters are fixed, based on experimental results. If necessary, these parameters can be changed manually. The remaining parameters are manipulated through the use of a configuration file. The configuration file is the only argument for the program. An example configuration file is provided in setup.dat.

Parameters listed in this file include: two input images, image masks, percent of edges to be detected, output file name, contrast enhancement flag (0=on, 1=off), metric choice (HD, GHD, or RGHD), radius for grayscale image creation, tolerance for GHD, maximum deformation expected (radius of search window) for GHD, radius of smoothing window, percentage for robust statistic, and minimum number of elements for smoothing window.

As well as creating a local distance map, `RunAssessment.cxx` outputs local distance estimation statistics to the standard out. These global statistics include maximum, root mean square, 90% percentile, mean, least trimmed squares, and least median squares.

5.5 DFGenerator.cxx

The synthetic Gaussian deformation used to obtain the results in the paper, as described in Section 2, is implemented in `DFGenerator.cxx`. Two versions of `DFGenerator` have been implemented (2D and 3D). An N-d implementation as an ITK class is left as a future work. The deformation vectors are selected using `itk::GaussianDistribution` and then interpolation is done using `itk::ThinPlateSplineKernelTransform`. The deformation field is controlled through the spacing and variance arguments. The program outputs all of the details necessary to recreate the deformation: the deformed image, deformation field, deformed mask, and deformation norm.

5.6 Evaluation.cxx

A simple evaluation tool is also included, `evaluation.cxx`. This tool is used to compare “ground truth” (i.e. deformation field norm) to local distance estimation. The program iterates through the two input images, comparing the difference in corresponding voxels, and calculates the percentage of outliers as defined in Section 2. In addition, this program outputs the percentage of outliers that are underestimates. An underestimate is defined as any voxel in the local distance estimation whose value is less than the “ground truth.”

5.7 Software Requirements

You need to have the following software installed:

- Insight Toolkit 3.4.0 (the version used to develop the software)
- CMake 2.4

This document was created using \LaTeX , with the graph Figures produced by [xmgrace](#), and diagrams created in [Kivio](#). Image data were visualized with ImageViewer from InsightApplications, and [Paraview](#).

6 Conclusions

We have presented a method for automated assessment of misalignment error. We have implemented 3 new ITK filters `itk::LocalGrayscaleDistanceMapImageFilter`, `itk::LocalGrayscaleDistanceMapImageFilter` and `itk::LocalDistanceMapSmoothingImageFilter` for use in this automated assessment and have provided the code for evaluation to reproduce the results presented in this paper. The results have shown RGHD is more robust in terms of outliers than other methods discussed and can potentially improve the accuracy of image alignment assessment. Furthermore, the local error estimation method we introduce has several applications in registration assessment. First, it can yield a global similarity metric which can be used for registration comparison or assessment of registration quality. Second, it can provide visual assessment of local error estimation as shown in Figure 8. Future versions of this work will include a templated version of the deformation generator in addition to further studies of parameter selection.

7 Acknowledgments

Experimental images have been provided by [Brainweb Simulated Brain Database](#) and [Insight Toolkit](#). This work was supported in part by NSF grants CSI-0719929 and CNS-0312980, and by the John Simon Guggenheim Memorial Foundation.

References

- [1] N. Archip, O. Clatz, S. Whalen, D. Kacher, A. Fedorov, A. Kot, N. Chrisochoides, F. Jolesz, A. Golby, P. Black, and S.K. Warfield. Non-rigid alignment of pre-operative MRI, fMRI, and DT-MRI with intra-operative MRI for enhanced visualization and navigation in image-guided neurosurgery. *Neuroimage*, 35:609–624, 2007. [1](#), [2](#), [4.1](#)
- [2] N. Archip, S. Tatli, P. Morrison, F. Jolesz, S.K. Warfield, and S. Silverman. Non-rigid registration of pre-procedural MR images with intra-procedural unenhanced CT images for improved targeting of tumors during liver radiofrequency ablations. *Proc. of MICCAI'07*, pages 969–977, 2007. [1](#), [2](#), [4.1](#)
- [3] B. Baudrier, F. Nicolier, G. Millon, and S. Ruan. Binary-image comparison with local-dissimilarity classification. *Pattern Recognition*, 41:1461–1478, April 2008. [1](#), [2](#), [2](#)
- [4] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8:679–714, 1986. [2](#)
- [5] C. Cocosco, V. Kollokian, R. Kwan, and A. Evans. Brainweb: Online interface to a 3d MRI simulated brain database. *NeuroImage*, 5(4), 1997. [4](#)
- [6] D. Huttenlocher, D. Klanderman, and A. Rucklidge. Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, September 1993. [1](#), [2](#)
- [7] R. Kwan, A. Evans, and G. Pike. MRI simulation-based evaluation of image-processing and classification methods. *IEEE Transactions on Medical Imaging*, 19(11):1085–97, November 1999. [4](#)
- [8] F. Morain-Nicolier, S. Lebonvallet, E. Baudrier, and S. Ruan. Hausdorff distance based 3d quantification of brain tumor evolution from MRI images. *Proc. of 29th Annual Intl Conf of the IEEE EMBS*, pages 5597–5600, 2007. [1](#)
- [9] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990. [2](#)
- [10] P. Rogelj, S. Kovačič, and J.C. Gee. Point similarity measures for non-rigid registration of multi-modal data. *Comput. Vis. Image Underst.*, 92(1):112–140, 2003. [3](#)
- [11] P. Rousseeuw and A. Leroy. *Robust Regression and Outlier Detection*. John Wiley and Sons, 1987. [2](#)
- [12] J. Stark. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transaction on Image Processing*, May 2000. [2](#)
- [13] C. Zhao, W. Shi, and Y. Deng. A new hausdorff distance for image matching. *Pattern Recognition Letters*, 26(5):581–586, April 2005. [1](#), [2](#)

The Use of Robust Local Hausdorff Distances in Accuracy Assessment for Image Alignment of Brain MRI

Release 0.20

Eric Billet, Andriy Fedorov and Nikos Chrisochoides

May 9, 2008

Center for Real Time Computing
Department of Computer Science
College of William and Mary
Williamsburg, VA 23185
{ejbill,fedorov,nikos}@cs.wm.edu

Abstract

We present and implement an error estimation protocol in the Insight Toolkit ([ITK](#)) for assessing the accuracy of image alignment. We base this error estimation on a robust version of the Hausdorff Distance (HD) metric applied to the recovered edges of the images. The robust modifications we introduce to the HD metric significantly reduce the amount of outliers in the local distance error estimation. We evaluate the accuracy of our protocol on synthetically deformed images. We provide the source code and datasets to reproduce this evaluation. The proposed method is shown to improve error assessment when it is compared with conventional HD methods. This approach has many applications including local estimation and visual assessment of registration error and registration parameter selection.

Contents

1	Introduction	2
2	Methods	3
3	Experimental Framework	4
4	Results	6
4.1	3D Image Results	6
4.2	2D Image Results	8
5	Implementation	9
5.1	itk::LocalDistanceMapImageFilter	10
5.2	itk::LocalGrayscaleDistanceMapImageFilter	10
5.3	itk::LocalDistanceMapSmoothingImageFilter	12
5.4	RunAssessment.cxx	12

5.5	DFGenerator.cxx	13
5.6	Evaluation.cxx	13
5.7	Software Requirements	13
6	Conclusions	13
7	Acknowledgments	14

1 Introduction

Assessing the accuracy of image alignment is a vital step in image registration. Accuracy assessment is essential for verification of registration quality, comparison of registration results, and registration parameter selection. However, estimating registration accuracy is a challenging problem for the medical imaging community. Clinical application of registration methods require the validation of registration results, but this validation can only be performed on cases where ground truth is known (cadavers, invasive markers, phantoms). The ability to assess registration quality of non-rigid brain registration during image guided neurosurgery is limited; an expert can assess quality by locating a series of clearly identifiable landmarks or regions in both intra-operative and pre-operative images. Arguably, this may provide the best precision, but the assessment is limited to the easily identifiable landmarks. This method is also time consuming and requires the involvement of an expert which may not be practical during neurosurgery, when comparison of thousands of registrations is necessary. Consider a scenario when the optimum parameters of a registration method are not known, and we need to perform a large scale parametric search for the best parameter configuration. In this case, we are facing comparison of hundreds and even thousands of registration results, which in practice is only feasible using automatic or semi-automated approaches. We attempt to automate this assessment process through the implementation of local similarity metrics.

Although, similarity metrics are no substitute for validation, similarity metrics have been used to assess registrations. They can provide confidence in the registration result as well as criteria for comparison between registration methods and parameters. *One common shortcoming of this approach is that similarity metric values do not provide quantification in terms of distance in physical space.* Thus, one has the ability to track improvements among registrations but cannot speculate about the degree of misalignment

One approach to this problem is the use of Hausdorff Distance (HD) metric. The HD metric is derived from Euclidean distances between point sets and thus related to local alignment distances unlike many other image similarity measures. The original HD metric is highly sensitive to noise [6], but variations of the metric have been proposed to improve the robustness of the metric. A survey of these modifications can be found in [13, 3]. Even with the correspondence to physical space, HD based metrics have scarcely been used for assessment of image alignment. They were used by Morain-Nicolier et. al [8] to quantify brain tumor evolution and later by Archip et al. [1, 2] to assess performance of non-rigid image registration of brain MRI. To our knowledge no formal study has been completed to investigate the use of HD as an accuracy assessment measure for image alignment. We evaluate existing metrics and implement an automated assessment approach based on a robust variation of the HD metric. We provide three ITK filters for local assessment and implement an evaluation framework for comparison of these methods. Based on our findings, our robust HD approach can significantly improve the results as compared to the previously used version of HD.

2 Methods

Problem: assess accuracy of alignment between two images, I and J .

We develop a method to solve this problem through estimation of point-wise alignment error. The HD metric operates on sets of points, and in order to apply it to the problem we need first to identify such (feature) points in the images. Consistency of feature points identification is essential to the robustness and correctness of the metric. Thus, prior to local estimation, several preprocessing steps must be employed on the input images.

First, we smooth the input images using anisotropic diffusion [9] to filter out noise. Anisotropic diffusion smoothing is chosen because it preserves features and all of the similarity metrics we examine are based on comparison of features of images rather than intensities. Second, adaptive contrast equalization [12] is performed to improve the edge detection results. Third, edges are identified with a Canny edge detection filter [4]. Thresholds for Canny edge detection are adaptively selected using binary search to have approximately the same number of edges in each image. The number of edge points is controlled by changing the upper and lower values of the threshold parameters in the edge detector [4]. All preprocessing steps were implemented using ITK filters. The implementation details can be found in Section 5.

Once the feature edges (points) are identified in both images, we can depart from the images and analyze the alignment accuracy by working with the two sets of points. Let A and B be the set of binary images created from extracting features from I and J , respectively, and $A = \{a_1, a_2, \dots\}$ and $B = \{b_1, b_2, \dots\}$ be the corresponding set of feature points (non-zero pixels). The directed HD between these two sets of points, $h(A, B)$, is defined as the maximum distance from any point in A to any point in B . The symmetric HD, $H(A, B)$ is the maximum of both direct distances [6]:

$$h(A, B) = \max_{a \in A} (d(a, B)), \text{ where } d(a, B) = \min_{b \in B} \|a - b\| \quad (1)$$

$$H(A, B) = \max(h(A, B), h(B, A)) \quad (2)$$

HD provides a global comparison of similarity, and was used by Archip et al. [1, 2] to estimate the alignment of brain MRI images after non-rigid registration. However, the notion of local similarity is lost. For example, HD values can remain constant whether a set of images are misaligned in just one region or many regions. The local distance map (LDM) proposed by Baudrier et al. [3] is more suitable for point-wise error estimation.

The LDM derives a local measure of dissimilarity for 2D binary images by comparing images locally with a sliding window, where $H(A, B)$ is used as a dissimilarity measure within the window.




Figure 1: Binary image (gray/white cells) with corresponding grayscale image (D=3) after non-rigid registration.

$$\forall x \in \mathbf{R}^3 : H_{loc}(x) = |\mathbf{1}_{\mathcal{A}(x)} - \mathbf{1}_{\mathcal{B}(x)}| \times \max(d(x, A), d(x, B)), \quad (3)$$

$$\mathbf{1}_{\mathcal{A}x} = \begin{cases} 1 & \text{if } \mathcal{A}(x) \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{where } \mathcal{A}(x) \text{ is the voxel value of } A \text{ at location } x.$$

Remark 1. The conventional $H(A, B)$ can be derived by taking the maximum of the local, H_{loc} , comparisons i.e., $H(A, B) = \max(H_{loc}(A, B))$.

LDMaP is parameterless and yields a more descriptive measure of *local* similarity than HD, but the LDMaP as described by Baudrier et al. [3] does not improve robustness of local estimation. Ideally, we are looking for a measure which is equal or analogous to the distance between corresponding points in the two images. However, this cannot be assumed for HD, since HD has no point-to-point correspondence.

In this paper we extend the definition of LDMaP by adding the notion of correspondence by using a grayscale modification of the HD. Grayscale HD was proposed by Zhao et al. [13] for matching 2D images corrupted by noise. We extend this grayscale HD definition to 3D for medical image comparison and include it within the LDMaP definition.

Given two input binary images, A and B , let \tilde{A} and \tilde{B} be the grayscale images whose voxel values are initialized to the number of non-zero voxels in constant size neighborhood of the corresponding binary image. \tilde{A} and \tilde{B} have the same size and boundary as the initial binary images. The neighborhood is isotropic and of size $D \times D \times D$. Figure 1 provides a 2D example with $D = 3$. This modification improves the sense of point correspondence while calculating point distances because the corresponding point in the second image is likely to have similar or same number of neighbor feature points. The directed distance $d(a_g, \tilde{B})$ where g is the grayscale value at voxel a is defined as the minimum distance from point a to any point in \tilde{B} with grayscale value within some tolerance, t , of g . The tolerance allows to adjust the sensitivity to differences in edge images (i.e., from resampling, different imaging device, etc.).

$$d(a_g, \tilde{B}) = \min_{b_{g'} \in \tilde{B}} \|a_g - b_{g'}\|, \quad g - t \leq g' \leq g + t \quad (4)$$

Using Equation 4 we define the local grayscale HD, GH_{loc} as follows:

$$\forall x \in \mathbf{R}^3 : GH_{loc}(x) = |\mathbf{1}_{\mathcal{A}}(x) - \mathbf{1}_{\mathcal{B}}(x)| \times \max(d(x_g, \tilde{A}), d(x_g, \tilde{B})). \quad (5)$$

Similar to the Hausdorff Distance, Grayscale Hausdorff Distance (GHD) $GH(A, B)$ is defined as the maximum of the local calculations, $GH(A, B) = \max(GH_{loc}(A, B))$.

In addition, we further improve the robustness of Grayscale Hausdorff Distance by applying robust statistics. Let $RGH_{loc}(x)$ be the local Robust Grayscale Hausdorff Distance of images A and B at voxel x , and let W_x be an isotropic window of size $S \times S \times S$ centered around x . $RGH_{loc}(x)$ is defined as the robust average calculated from the voxels inside W_x . Least trimmed squares [11] is a reasonable choice for this robust average, but other averages can be employed to improve robustness. We define the Robust Grayscale Hausdorff Distance (RGHD) $RGH(A, B) = \max(RGH_{loc}(A, B))$.

The complete diagram of the alignment accuracy assessment framework, together with the parameters used at each step, is depicted in Figure 2.

3 Experimental Framework

We evaluate the effectiveness of the proposed accuracy assessment methodology for non-rigid registration of brain MRI using the synthetic ground truth data. First we use the method described by Rogelj et al. [10] to construct a synthetic deformation field. The synthetic deformation field is applied to the original grayscale image, followed by feature detection step performed on both the original and deformed images. The proposed error recovery methodology is then employed to estimate the misalignment between the original and

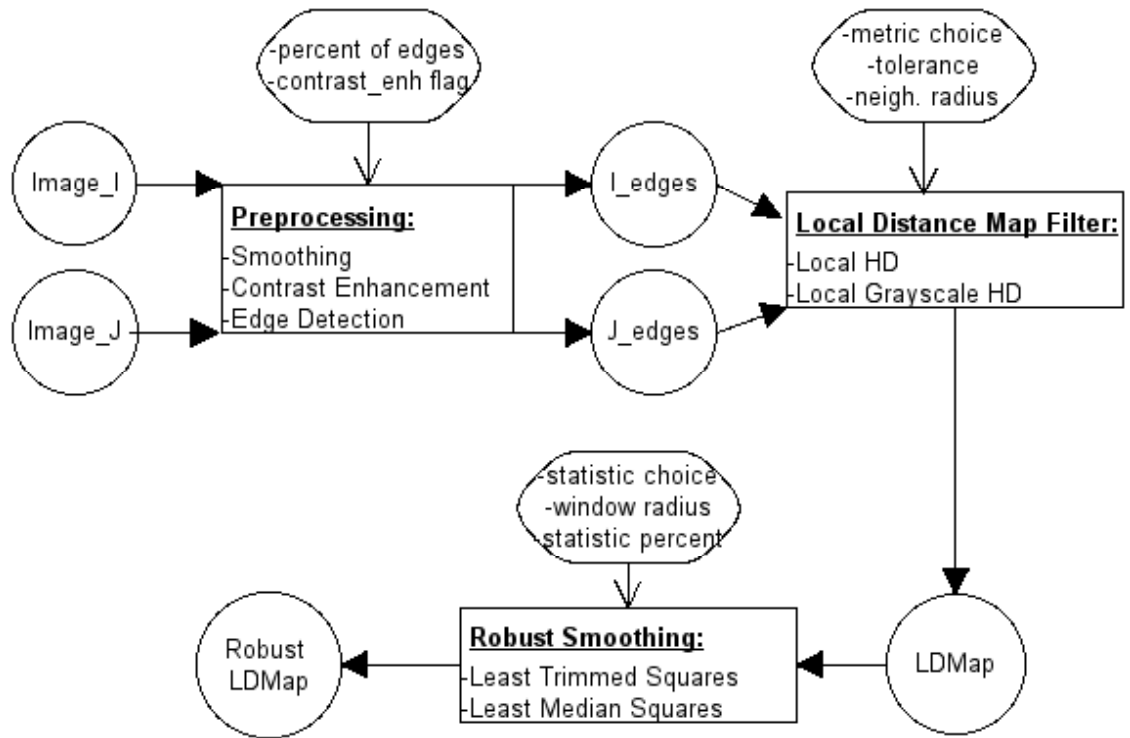


Figure 2: The basic components of the assessment framework: (1) preprocessing of input images, (2) computation of local distance metric on edge images and (3) robust smoothing of local distance map.




Figure 3: Random deformation vectors are generated at the knots (red circles) of the deformation field grid that are located within the user-defined mask.

the deformed images at the selected feature points. The true misalignment value, which is the magnitude of the synthetic deformation at a point, ideally should be equal to the error value recovered by the either of the proposed assessment measures.

In order to create the synthetic deformation field, we first construct a sparse point sample at the knots of an isotropic sampling grid. This procedure is illustrated in Figure 3, where point sample is shown with red circles. We assign a random deformation vector at each of these points, with the components of the vector drawn from a Gaussian distribution parameterized by mean($\mu = 0$) and variance(σ). The dense deformation field is constructed by using Thin Plate Splines interpolation at non-knot image points within the binary mask.

The current implementation of the synthetic deformation is not provided as a generic class. Instead, we have separate tools to construct 2D and 3D deformation fields. The deformation fields are 0 by the sampling grid spacing and the variance of the Gaussian distribution, which should be provided by the user, together with the binary mask of the region of interest. `itk::ThinPlateSplineKernelTransform` is used to interpolate the values of the dense deformation field at non-grid pixels

We compare the local estimation methods with ground truth error using two measures: distribution of error and percentage of outliers. Ideally, the distribution of local error estimates (H_{loc} , GH_{loc} and RGH_{loc}) will closely mimic the true error distribution. Also, a good estimate of error should have minimum number of outliers. Let d_i be the local distance at point location i measured by a local estimation method and e_i be the true error at the same point. We define outliers as any point i where $|d_i - e_i| > 2mm$. We choose 2 mm because the deformation field is in physical space and the HD distance implementation is limited to 1 mm image spacing. Thus, errors as large as $\sqrt{3}$ cannot be prevented.

4 Results

We present results for each of the described metrics (HD, GHD, RGHD) using the evaluation framework described in Section 3 for both 3D and 2D images. The 3D images are not provided with this submission due to possible copyright issues, but can be freely downloaded from BrainWeb Simulated Brain Database [5, 7]. We used the following settings for constructing the simulated image: Modality=T1, Slice thickness = 1mm, Noise = {0%,9%} and RF=0%. The 2d images accompanying this paper were taken from the example data provided with ITK (ITK/Examples/Data/BrainProtonDensitySliceBorder20.png, BrainProtonDensitySliceBorder20Mask.png).

4.1 3D Image Results

The 3D results presented in this paper are produced using the following parameters. The variance parameter, σ , of the Synthetic Gaussian deformation was set to the integer values between 1 and 12 to get increasingly more complex deformations. The deformation grid spacing was set at 30 mm for all test cases. These settings create displacements with a maximum of 12 mm and an average between 1 and 4 mm, depending on variance. Parameters for GHD include: neighborhood size, for creation of grayscale images, set to $D = 3$ and tolerance for $GHD_{loc}(x)$, $t = 2$. RGHD used a least trimmed squares robust metric with percentage set to 80% and window size, $W_x = 11$. We are able to incorporate the use of noisy images in our evaluation because Brainweb provides a simulation of noisy images, and thus we evaluate local alignment assessment between undeformed images with 0% noise and deformed images with 0% and 9% noise.

The error distributions for: the actual error, H_{loc} , GH_{loc} , and RGH_{loc} for $\sigma = 5$ of the synthetic deformation




Figure 4: Distribution of the error and the HD, GH, and RGH values for the same synthetic deformation case (3D BrainWeb image, Gaussian deformation, variance 5).




Figure 5: The change in outlier percentages depending on the variance of Gaussian Deformation. Left: deformed 2D image included with paper Right: deformed 3D Brainweb images with 0% and 9% noise.




Figure 6: Global error statistics (i.e. HD, 95%HD, RGHD) for synthetically deformed 3D BrainWeb images with and without noise. RGHD results were derived from RGH_{loc} values.

are shown in Figure 4. Other variances had similar results. The distribution of RGH_{loc} is an improved approximation of the error distribution over H_{loc} and GH_{loc} .

As expected the data in Figure 5 indicates that RGH_{loc} is also the most robust method. Figure 5 displays the percentage of outliers in the set of local distance estimations. RGH_{loc} remains at a low percentage of outliers throughout the 3D testing. GH_{loc} also provides improved robustness over the H_{loc} estimations. For best results, we recommend the use of RGH_{loc} when deformations of magnitude 3mm to 7mm are present.

The error distribution and outliers help to build confidence in the robustness and accuracy of the local estimates. Next we examine the use of robust local estimation methods for non-rigid registration of brain MRI comparison. Our hypothesis is that summary statistics of local estimations provide a more robust global statistic for comparison purposes. We test this hypothesis by examining the 95% percentile of the local estimates at each variance and compare it with the results obtained with the conventional HD and the 95% PHD used by Archip et al. [1, 2]. These results are presented in Figure 6. RGHD tracks the improvement well (i.e. increasing constantly w.r.t. error) even in the presence of noise. Moreover, RGHD values consistently over-estimate (upper-bound) the true mean error. This is expected because the robust metrics are based on the assumption that the majority of outliers are underestimates and thus the calculated mean increases. The use of other summary statistics (RMS, LTS) on the local estimates may improve this discrepancy, but the study of these statistics is left to future work.

4.2 2D Image Results

We evaluate alignment assessment between the undeformed and deformed 2D ITK images with the following parameters. The deformation field is created with: $\sigma = \{1, 2, 3, \dots, 12\}$ and spacing = 30mm. Deformation vectors with magnitude averaging between 1.5mm to 7mm are observed, magnitudes increase with variance. A maximum deformation of 9mm occurs at $\sigma = 12$. The parameters for local estimation parameters are based on the 3D results: neighborhood of GH_{loc} , $D = 3$, tolerance of grayscale, $t = 1$, window size of RGH_{loc} , $W_x = 11$, and percentage = 80%. Adjustments of these parameters may cause different results. A more exhaustive study of parameter selection will be performed in the future and reported in upcoming updates of this document.

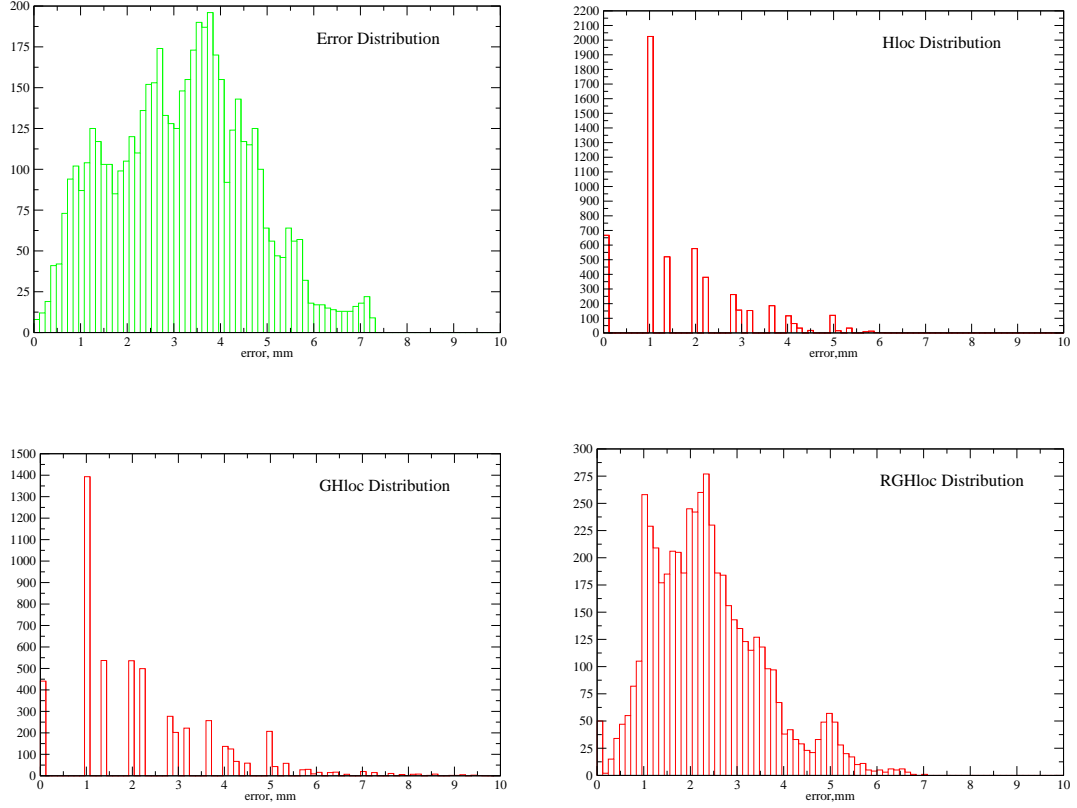


Figure 7: Distribution of the error and the HD, GH, and RGHD values for the same synthetic deformation case (2D ITK image, Gaussian deformation, variance 8).

Similar to the results in 3D data, RGHD provides the closest approximation to true error distribution with 2D images, see Figure 7. However, as seen in Figure 5 the outliers for RGHD are not as low with the 2D images. Many factors could contribute to this difference including the parameters chosen and the deformation field itself. Although, the variances were set to the same values, the deformation magnitudes are slightly higher on average in 2D image tests. The RGHD is still shown to be the most robust method compared.

We also examined the use of local estimation for visualizing the error in Figure 8. We color the LDMMap created by RGH_{loc} estimations, allowing visualization of the magnitude of error throughout the image.

5 Implementation

The presented local error estimation methods have been implemented in three n-dimensional image to image filters: `itk::LocalDistanceMapImageFilter`, `itk::LocalGrayscaleDistanceMapImageFilter`, and `itk::LocalDistanceMapSmoothingImageFilter`. In addition, the entire alignment assessment process is implemented in `RunAssessment.cxx`. In Figure 5 we present the entire framework for our evaluation of local assessment methods. Descriptions of each of these filters and their use is described in the following sections. Two helper classes were also implemented. `itk::CountImageFilter` assists in the conversion of a binary image to its grayscale counterpart by implementing a simple counter of non-zero voxels in a predefined neighbor-




Figure 8: Local estimation of misalignment using RGHD. Left: Undeformed 2D image. Center: Deformed image, Gaussian variance of 8 mm. Right: LDM of the deformed and undeformed images using RGHD, color-coded, mm.

hood region. `itk::errorStatistics` is used to calculate descriptive error statistics from the robust LDM image, see Figure 2.

5.1 itk::LocalDistanceMapImageFilter

`itkLocalDistanceMapImageFilter` is the ITK implementation of local HD metric as described in Section 2. This filter is derived from the `itk::ImageToImage` filter and is templated over the two input images and the output image. This filter requires two distance maps as input. It outputs a local distance map with values of each voxel, x , equal to $H_{loc}(x)$.

This processing step does not have any parameters, and the only methods of relevance are those that provide the set and get functionality for the filter inputs:

- `SetInput1(const TDistance1 *)`, `GetInput1()`
- `SetInput2(const TDistance1 *)`, `GetInput2()`

The class is multi-threaded and the entire algorithm is contained in `ThreadedGenerateData()`.

5.2 itk::LocalGrayscaleDistanceMapImageFilter

We implemented a local grayscale Hausdorff filter based on the definition in Section 2. Similar to the previous filter, `itk::LocalGrayscaleDistanceMapImageFilter` is derived from the `itk::ImageToImage` filter and is templated over its two input images and one output image. This filter takes two binary edge images as input and outputs a local distance map with values of each voxel, x , equal to $GH_{loc}(x)$.

The first step is the construction of the grayscale image from the input binary image, as shown in Figure 1. This step is facilitated by the helper class `itk::CountImageFilter` we implemented, and is 0 by the neighborhood radius value. As described in Section 2, to find the $GH_{loc}(x)$ values we must perform a search at each feature point, for the nearest neighbor features with values within a given range, $g - t \leq g' \leq g + t$. Exhaustive search is computationally expensive, instead the $GH_{loc}(x)$ values are computed by performing an

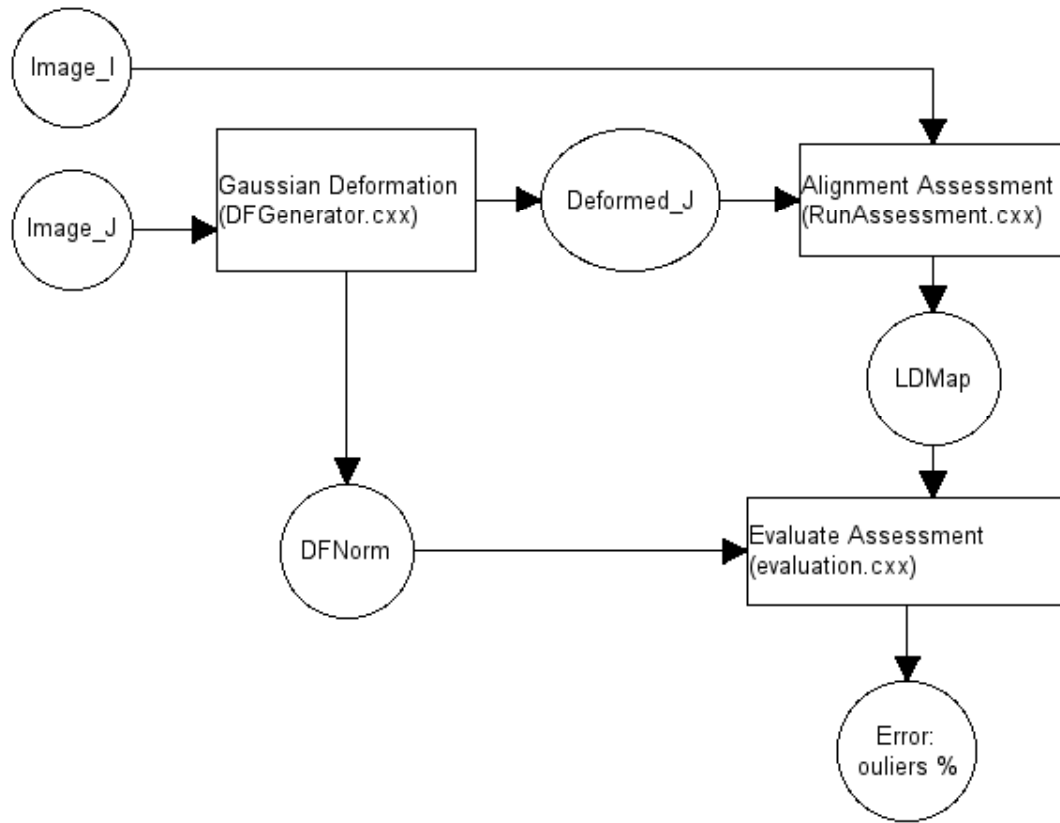


Figure 9: Components for evaluation of accuracy alignment.

expanding search with maximum radius, `MaxDef`, which is another parameter of the filter. `MaxDef` should be chosen based on the estimated maximum deformation value expected. The indices for the expanding search are determined in `BeforeThreadedGenerateData()`.

The input methods for this filter are:

- `SetInput1(const TEdge1 *)`, `GetInput2()`
- `SetInput2(const TEdge1 *)`, `GetInput2()`

The following parameters for this filter can be manipulated through `Set/Get` functions:

- **MaxDef**: This controls the search radius for finding the nearest neighbor, should be set to the maximum deformation possible (e.g. 10% of adult brain size). If no feature point is found within the search radius, the voxel is set to an unrealistic value (-100).
- **Tol**: the tolerance, t , for the local grayscale Hausdorff metric. Higher values of t will create smaller distances, but the 0 of corresponding points will decrease.
- **Radius**: radius of the neighborhood size, D , ($2 \times \text{radius} + 1 = D$), for which the grayscale images are created. Typical values used are 1 to 4.

The class is multi-threaded. The generation of grayscale images is performed in `BeforeThreadedGenerateData()`, but the rest of the calculations for local Grayscale HD is contained in `ThreadedGenerateData()`.

5.3 itk::LocalDistanceMapSmoothingImageFilter

itk::LocalDistanceMapSmoothingImageFilter is a smoothing filter designed specifically for post-processing of the local distance map to obtain the values of RGH_{loc} as described in Section 2. This filter applies a robust average within a window centered around a feature point. The filter only takes into account feature points. It accomplishes this by using “edge” (feature point) images as a mask during computation. The local robust grayscale Hausdorff distances (RGH_{loc}) are computed by applying this filter to the local distance map created by itk::LocalGrayscaleDistanceMapImageFilter.

The input methods for this filter are:

- SetInput1(const TEdge1 *), GetInput1() the first edge image, typically used to create local distance map
- SetInput2(const TEdge2 *), GetInput2() the second edge image, typically used to create local distance map
- SetInput3(const TDistance1 *), GetInput3() the local distance map

The parameters for this filter are manipulated by Get/Set commands:

- **Radius:** the radius, $(2 \times radius + 1 = S)$, of the window used for robust averaging
- **MinElements:** the minimum number of feature points needed to compute average. If below minimum assume there is no confidence in region and voxel set to unrealistic value (-100).
- **Percent:** the percentage used for the robust statistic, (i.e. $|1 - percent|$ is discarded in least trimmed squares)

Function GetRobustStat(std::vector) computes the robust statistic on the distances inside the window. Currently, only LTS is implemented, but other robust statistics could easily be added. As with the other filters, this filter is multi-threaded and derived from the ImageToImage filter.

5.4 RunAssessment.cxx

An example code for running the complete alignment assessment process on two images is provided in RunAssessment.cxx. Three basic components of the process (preprocessing, local distance map generation, robust smoothing) are each implemented. The preprocessing is performed using three ITK filters: itk::CannyEdgeDetectionImageFilter, itk::CurvatureAnisotropicDiffusionImageFilter, and itk::AdaptiveHistogramEqualizationImageFilter. Smoothing and contrast enhancement parameters are fixed, based on experimental results. If necessary, these parameters can be changed manually. The remaining parameters are manipulated through the use of a configuration file. The configuration file is the only argument for the program. An example configuration file is provided in setup.dat.

Parameters listed in this file include: two input images, image masks, percent of edges to be detected, output file name, contrast enhancement flag (0=on, 1=off), metric choice (HD, GHD, or RGHD), radius for grayscale image creation, tolerance for GHD, maximum deformation expected (radius of search window) for GHD, radius of smoothing window, percentage for robust statistic, and minimum number of elements for smoothing window.

As well as creating a local distance map, `RunAssessment.cxx` outputs local distance estimation statistics to the standard out. These global statistics include maximum, root mean square, 90% percentile, mean, least trimmed squares, and least median squares.

5.5 DFGenerator.cxx

The synthetic Gaussian deformation used to obtain the results in the paper, as described in Section 2, is implemented in `DFGenerator.cxx`. Two versions of `DFGenerator` have been implemented (2D and 3D). An N-d implementation as an ITK class is left as a future work. The deformation vectors are selected using `itk::GaussianDistribution` and then interpolation is done using `itk::ThinPlateSplineKernelTransform`. The deformation field is controlled through the spacing and variance arguments. The program outputs all of the details necessary to recreate the deformation: the deformed image, deformation field, deformed mask, and deformation norm.

5.6 Evaluation.cxx

A simple evaluation tool is also included, `evaluation.cxx`. This tool is used to compare “ground truth” (i.e. deformation field norm) to local distance estimation. The program iterates through the two input images, comparing the difference in corresponding voxels, and calculates the percentage of outliers as defined in Section 2. In addition, this program outputs the percentage of outliers that are underestimates. An underestimate is defined as any voxel in the local distance estimation whose value is less than the “ground truth.”

5.7 Software Requirements

You need to have the following software installed:

- Insight Toolkit 3.4.0 (the version used to develop the software)
- CMake 2.4

This document was created using \LaTeX , with the graph Figures produced by [xmgrace](#), and diagrams created in [Kivio](#). Image data were visualized with ImageViewer from InsightApplications, and [Paraview](#).

6 Conclusions

We have presented a method for automated assessment of misalignment error. We have implemented 3 new ITK filters `itk::LocalGrayscaleDistanceMapImageFilter`, `itk::LocalGrayscaleDistanceMapImageFilter` and `itk::LocalDistanceMapSmoothingImageFilter` for use in this automated assessment and have provided the code for evaluation to reproduce the results presented in this paper. The results have shown RGHD is more robust in terms of outliers than other methods discussed and can potentially improve the accuracy of image alignment assessment. Furthermore, the local error estimation method we introduce has several applications in registration assessment. First, it can yield a global similarity metric which can be used for registration comparison or assessment of registration quality. Second, it can provide visual assessment of local error estimation as shown in Figure 8. Future versions of this work will include a templated version of the deformation generator in addition to further studies of parameter selection.

7 Acknowledgments

Experimental images have been provided by [Brainweb Simulated Brain Database](#) and [Insight Toolkit](#). This work was supported in part by NSF grants CSI-0719929 and CNS-0312980, by the John Simon Guggenheim Memorial Foundation, and by the Commonwealth of Virginia under the project name “Building Research Capacity in Biomedical Research and Biomaterials Engineering.”

References

- [1] N. Archip, O. Clatz, S. Whalen, D. Kacher, A. Fedorov, A. Kot, N. Chrisochoides, F. Jolesz, A. Golby, P. Black, and S.K. Warfield. Non-rigid alignment of pre-operative MRI, fMRI, and DT-MRI with intra-operative MRI for enhanced visualization and navigation in image-guided neurosurgery. *Neuroimage*, 35:609–624, 2007. [1](#), [2](#), [4.1](#)
- [2] N. Archip, S. Tatli, P. Morrison, F. Jolesz, S.K. Warfield, and S. Silverman. Non-rigid registration of pre-procedural MR images with intra-procedural unenhanced CT images for improved targeting of tumors during liver radiofrequency ablations. *Proc. of MICCAI’07*, pages 969–977, 2007. [1](#), [2](#), [4.1](#)
- [3] B. Baudrier, F. Nicolier, G. Millon, and S. Ruan. Binary-image comparison with local-dissimilarity classification. *Pattern Recognition*, 41:1461–1478, April 2008. [1](#), [2](#), [2](#)
- [4] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8:679–714, 1986. [2](#)
- [5] C. Cocosco, V. Kollokian, R. Kwan, and A. Evans. Brainweb: Online interface to a 3d MRI simulated brain database. *NeuroImage*, 5(4), 1997. [4](#)
- [6] D. Huttenlocher, D. Klanderman, and A. Rucklidge. Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, September 1993. [1](#), [2](#)
- [7] R. Kwan, A. Evans, and G. Pike. MRI simulation-based evaluation of image-processing and classification methods. *IEEE Transactions on Medical Imaging*, 19(11):1085–97, November 1999. [4](#)
- [8] F. Morain-Nicolier, S. Lebonvallet, E. Baudrier, and S. Ruan. Hausdorff distance based 3d quantification of brain tumor evolution from MRI images. *Proc. of 29th Annual Intl Conf of the IEEE EMBS*, pages 5597–5600, 2007. [1](#)
- [9] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990. [2](#)
- [10] P. Rogelj, S. Kovačič, and J.C. Gee. Point similarity measures for non-rigid registration of multi-modal data. *Comput. Vis. Image Underst.*, 92(1):112–140, 2003. [3](#)
- [11] P. Rousseeuw and A. Leroy. *Robust Regression and Outlier Detection*. John Wiley and Sons, 1987. [2](#)
- [12] J. Stark. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transaction on Image Processing*, May 2000. [2](#)
- [13] C. Zhao, W. Shi, and Y. Deng. A new hausdorff distance for image matching. *Pattern Recognition Letters*, 26(5):581–586, April 2005. [1](#), [2](#)