

TrueD Deformable Registration

Algorithm Description and Results Interpretation

www.usa.siemens.com/medical

SIEMENS
medical

1 Introduction

The clinical software application TrueD supports the reading of PET-CT scans of the same patient taken at up to three time-points (longitudinal follow-up), for example to allow the tracking of lesions over time and the monitoring of tumor response to therapy. In order to make comparisons between time-points easier, the software allows the CT datasets to be co-registered, providing a correspondence between points in the baseline study and the follow-up study (or studies). First of all, a rigid transformation is computed between pairs of datasets, allowing global rotation and translation of one dataset to align it with the other¹; following this rigid registration, it is then possible to perform a non-rigid (or deformable) registration, which allows local deformations of the study, to account, for example, for different patient positioning between scans or different levels of breathing. It is not possible for a rigid transformation alone to be able to account for these local differences throughout the study.

As an example, consider the study shown in Figure 1, with the top images showing rigid alignment of the two datasets, and the bottom showing the result of a deformable registration. The rigid alignment has provided gross alignment of all body organs and structures as a whole; organs are approximately in the same place when overlaying one image on the other, but mis-alignments due to respiration levels or patient positioning are visible (see around the ribs, the heart, and the bones of the arms in the coronal view). Deformable registration attempts to improve the alignment between corresponding locations in these regions: the ribs are aligned more closely and the arms are in a more similar position (note, however, that some differences may still remain if the shape or patient position is too dramatically different).




Figure 1: co-registration of two CT scans of the same patient (baseline in greyscale, follow-up overlaid in color). Top: rigid only; bottom: deformable. Note how the edges of the lung and the ribs are much better aligned with the deformable result.

2 Algorithm Details

Internally within the software, the deformable registration algorithm is based on the optical flow method of Lucas and Kanade [LuKa81], which solves for the displacement required to align each voxel in the scan using a brightness constancy constraint. This algorithm, originally designed to track objects in movie video sequences, has been specifically adapted to the purpose of this clinical application.

The brightness consistency constraint is built on the assumption that the intensity of any voxel in the source scan will be the same in the target scan, an assumption that is typically true in the case of CT datasets where Hounsfield units are used as standard. Mathematically, if we consider the image I such that $I(x, y, z, t)$ is the intensity of the baseline image at the point (x, y, z) , and $I(x + \delta x, y + \delta y, z + \delta z, t + \delta t)$ is the corresponding intensity in the follow-up image, then assuming that the displacements $(\delta x, \delta y, \delta z)$ required to align the two images are small², a Taylor series expansion can be used to give

$$I(x + \delta x, y + \delta y, z + \delta z, t + \delta t) \approx I(x, y, z, t) + \frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \frac{\partial I}{\partial z} \delta z + \frac{\partial I}{\partial t} \delta t \quad (2.1)$$

(neglecting higher order terms). Using the brightness constancy constraint

$$I(x + \delta x, y + \delta y, z + \delta z, t + \delta t) \approx I(x, y, z, t),$$

we obtain

$$\frac{\partial I}{\partial x} \delta x + \frac{\partial I}{\partial y} \delta y + \frac{\partial I}{\partial z} \delta z + \frac{\partial I}{\partial t} \delta t = 0$$

Dividing through by δt and writing partial derivatives as subscripts gives

$$I_x V_x + I_y V_y + I_z V_z = -I_t$$

where \vec{v} is the vector of displacements at the point (x, y, z) . In order to find the three unknowns, Lucas and Kanade suggest assuming a locally constant flow within a small window of the image of size $m \times m \times m$, for $m > 1$; this assumption leads to an over-determined system of $m^3 = n$ equations

$$\begin{pmatrix} I_{x_1} & I_{y_1} & I_{z_1} \\ I_{x_2} & I_{y_2} & I_{z_2} \\ \vdots & \vdots & \vdots \\ I_{x_n} & I_{y_n} & I_{z_n} \end{pmatrix} \vec{v} = \begin{pmatrix} -I_{t_1} \\ -I_{t_2} \\ \vdots \\ -I_{t_n} \end{pmatrix}$$

which can then be solved for \vec{v} using standard least squares techniques [GoVL].

2.1 Large Displacements

The algorithm presented in the previous section assumes that the displacements required to align the two images are small enough to be able to neglect the higher order terms in the Taylor series expansion (2.1), something that is unlikely to be the case for scans where deformable registration is required to make a significant improvement of the alignment. This situation can be dealt with using a standard technique, that of taking a hierarchical approach to the problem (see, e.g., [BeAn92]). Before proceeding, it is important to note that the sizes of the displacements relevant to optical flow are measured in terms of the voxel sizes of the image.

With hierarchical techniques, both baseline image I_b and follow-up image I_f are resampled to produce successively smaller versions (see Figure 2); the smallest images will only contain the low frequency (i.e., large feature) information present in the original images.

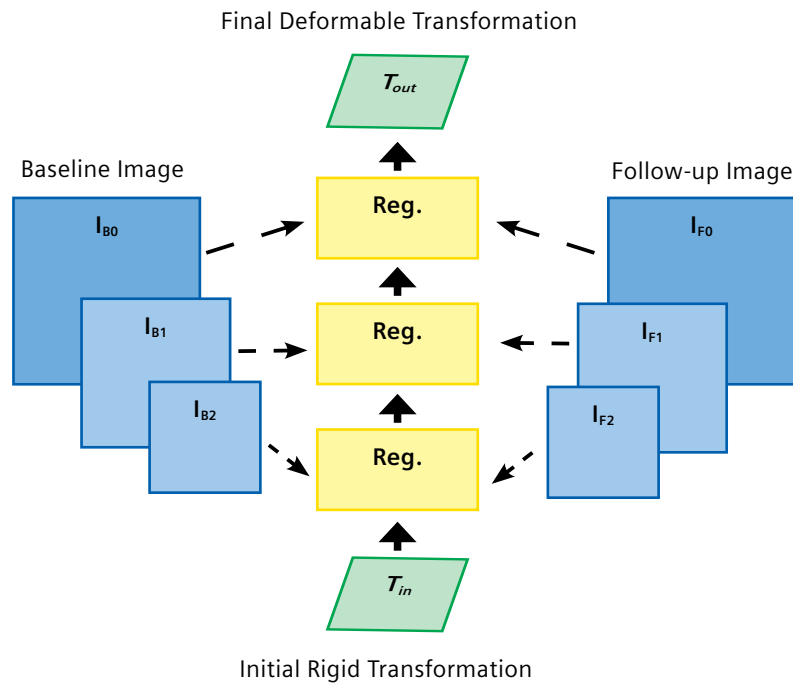


Figure 2: hierarchical computation of optical flow; starting with the initial rigid registration result, the two smallest images are registered to recover the large displacements; this transformation is then used to initialize the next levels of registration where successively finer levels of detail are recovered.

The key advantage of these small images is that displacements that may be large in the original image become relatively smaller in the resampled images since the voxels are larger, and the small displacement assumption is still valid. Starting by registering the smallest image pair, the resulting transformation can then be used to initialize the next level and so on, maintaining the assumption of small displacements throughout the algorithm, leaving registration of the larger images to recover the finest details.

2.2 Pre-processing Data

Before applying the deformable registration algorithm, the software internally applies a window of 350 and level of 40 to both baseline and follow-up scans to highlight features that the algorithm can use to match the images. This window/level gives good contrast between bones, air and tissue, as well as providing some level of soft-tissue contrast, helping to distinguish tissue-types that the algorithm will use when matching. However, it is important to note that for this window/level the bed will often appear saturated as one large, regular structure with no contrast, and hence it is possible, although unlikely, that unexpectedly large displacements can occur in this region of the scan.

3 Usage of Deformable Registration

Deformable registration is intended as a refinement of the rigid registration result, and hence it should be ensured that a good quality rigid registration is computed before performing a deformable registration. Failure to do this increases the likelihood of the deformable registration failing to produce clinically acceptable results.

3.1 Non-injective Transformations

With the registration algorithm as set out above, a correspondence between the baseline and follow-up images is computed such that for any given point in the baseline image, the corresponding point in the follow-up image can be computed. The converse, however, need not be true; for example, consider a case where a particular feature in the image has disappeared between scans – there may be several points within the baseline scan that all map to the same location in the follow-up scan. In this case, given that point in the follow-up scan, there is no single point in the baseline scan that directly corresponds, and the best that can be done is to choose one plausible solution.

This situation is fortunately very unlikely, but should serve as a warning that results should be carefully inspected where significant differences between the two scans exist, such as where surgery has been performed. Note that for typical usage, exact correspondence of all structures is not necessary: only the regions where uptakes are to be reported need match sufficiently well.

3.2 Accuracy of Results

Unfortunately, it is not possible to assert that deformable registration will always give valid results over the entire dataset (indeed, rigid registration suffers from similar, if not more severe limitations). In this respect, deformable registration is not dissimilar to Computer Assisted Detection (CAD) techniques: these techniques will have an operating point of sensitivity and specificity when applied to a particular problem; registration will similarly be shown to perform differently depending on how critical and accurate the results are expected to be, and how similar the images to be registered are (same positioning protocol, same breathing protocol, no dramatic changes in the patient anatomy between the two time points, etc.).

We have performed an internal study of the deformable registration algorithm in which a clinical expert compared the quality of the performance of the registration on a variety of cases (27 classic PET-CT follow-up studies with various diseases presented). It was recorded that the alignment was acceptable in 17 cases out of the 27 (for instance, providing results as displayed in Figure 1), but that some organs were misaligned in 10 cases. However, in 8 of these 10 cases, the areas where alignment failed were of no consequence to the assessment; in total, 25 out of 27 cases were considered acceptable for the assessment.

As with any registration, it is vital that the deformable registration result is carefully reviewed using the tools described in Section 3.3 in order to ensure that valid conclusions are drawn from the reading of the scan.

3.3 Assessment Of Results

In order to verify that deformable registration has computed a valid result, several tools are provided by the software to help the user in his or her assessment:

- linked crosshairs between baseline and follow-up images (Figure 3);
- displacement grid, showing regions of the image that have been stretched, compressed, or sheared (Figure 4, left);
- displacement map, showing relative amounts of deformation to different parts of the image (Figure 4, right).

Linked crosshairs (Figure 3) can be used for two purposes: firstly, they can be used to help assess the registration by checking the correspondence between identifiable physical structures such as bones or easily identifiable anatomical landmarks, and secondly, assuming that the registration has been verified, they can help identify correspondences between other structures (e.g., tumors).




Figure 3: linked crosshairs between two datasets show where the corresponding point from one dataset is in the other. Results should be carefully inspected to ensure that physical structures (e.g., bones) are correctly matched.

The displacement grid (Figure 4, left) provides a quick overview of regions of the image that have undergone stretching, compression, or shearing. If no deformation has occurred, regular squares will be seen; compression will produce squashed rectangles, whilst shearing will result in diamond-shapes. Special attention should be given to assessing the results in areas of high deformation (i.e., where the grid is not square).




Figure 4: the displacement grid (left) shows whether the deformation causes local compression, stretching, or shearing; the displacement map (right) gives an overview of the magnitude of displacements.

The displacement map (Figure 4, right) indicates the magnitude of the deformation at each point (note that this does not include any rigid transformation that may have been required to align the datasets initially). Cold colors mean little displacement and hot colors mean large displacements. As with the displacement grid, areas of high deformation should be inspected closely to verify the result.

4 Conclusion

In summary, the deformable registration provided by TrueD is a useful tool for improving local alignment of many longitudinal studies by registering the CT datasets following a rigid registration. However, as with all registration, the result should never be taken at face value, and a careful review should be performed using the assessment tools provided before drawing clinical conclusions.

References

- [BeAn92] Bergen, J.R., Anandan, P., Hanna, K.J., Hingorani, R., "Hierarchical Model-Based Motion Estimation", ECCV92: 237-252. 1992.
- [GoVL] Golub, G.H., Van Loan, C.F., "Matrix Computations", Third Edition, The Johns Hopkins University Press, 1996.
- [LuKa81] Lucas, B. D., Kanade, T., "An Iterative Image Registration Technique with an Application to Stereo Vision", Proceedings of 7th International Conference on Artificial Intelligence, 121-130. 1981.

About the author

Dr. Thomas Wright received a BA(hons) in Computation from the University of Oxford in 1999, and then completed a DPhil in Numerical Analysis (also at Oxford) in 2002. He has worked in the Science and Technology Team of Siemens Molecular Imaging (formerly Mirada Solutions) since the beginning of 2003.

On account of certain regional limitations of sales rights and service availability, we cannot guarantee that all products included in this brochure are available through the Siemens sales organization worldwide. Availability and packaging may vary by country and is subject to change without prior notice. Some/All of the features and products described herein may not be available in the United States.

The information in this document contains general technical descriptions of specifications and options as well as standard and optional features which do not always have to be present in individual cases.

Siemens reserves the right to modify the design, packaging, specifications and options described herein without prior notice. Please contact your local Siemens sales representative for the most current information.

Note: Any technical data contained in this document may vary within defined tolerances. Original images always lose a certain amount of detail when reproduced.

© 2007 Siemens Medical Solutions USA, Inc.
Order No. A9135-861-C1-4A00
Printed in USA 10-2007
All rights reserved

Contact Addresses

Siemens Medical Solutions USA
Molecular Imaging
2501 N. Barrington Road
Hoffman Estates, IL 60192-5203
USA
Telephone: +1-888-826-9702
www.siemens.com/mi

Siemens Medical Solutions USA
Molecular Imaging
810 Innovation Drive
Knoxville, TN 37932-2751
USA
Telephone: +1-888-826-9702
www.siemens.com/mi

Siemens Medical Solutions USA
Molecular Imaging
810 Innovation Drive
Knoxville, TN 37932-2751
USA

Headquarters
Siemens Medical Solutions USA
51 Valley Stream Parkway
Malvern, PA 19355-1406
USA
Telephone: +1-888-826-9702