

drop - User Guide

Version 1.05

Ben Glocker

September 24, 2009

Contents

1	Introduction	2
2	Files and Data	2
2.1	Loading Images	2
2.2	Data Masks	3
2.3	Saving Results	3
2.4	Applying the Deformation Field	4
2.5	Landmark-based Pre-Registration	4
3	Registration Parameters	4
3.1	Control Grid Settings	4
3.2	Label Set Settings	6
3.3	Optimization Settings	6
3.4	Parameter File	9

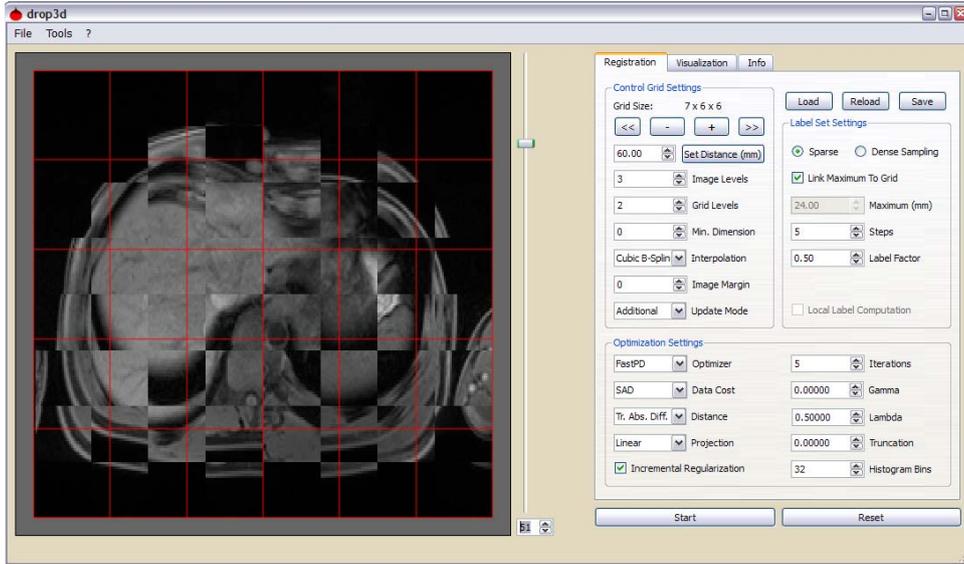


Figure 1: Overview of the graphical user interface.

1 Introduction

Drop is a software for deformable image registration using discrete optimization. The method is based on the publications [3, 4, 2]. In this user guide we describe the basic functionality of the software and the algorithm’s parameters. Hopefully, this document will help users to run our software in order to obtain good registration results. As this software presents research work in progress, bugs and typos of any kind are very likely, within this document as well in the software. We appreciate any comments, suggestions, or critics regarding our software. Feel free to contact us via mail: glocker@in.tum.de (Ben Glocker). Latest news, bug fixes, and new versions will always be announced on our website <http://www.mrf-registration.net>.

2 Files and Data

2.1 Loading Images

Images can be loaded using the *File* menu. Currently, the MHD format and the Analyze format are supported in both the 2D and 3D application. A simple DICOM support based on the GDCM library is provided as well. In 3D, it is possible to load single DICOM series located in one folder. In 2D it is also possible to load single DICOM slices as well as BMP, PNG, and PGM. Please note, that our registration is using gray value information only. RGB images will be converted to grayscale according to

$$I_{\text{gray}} = (11 \cdot I_{\text{red}} + 16 \cdot I_{\text{green}} + 5 \cdot I_{\text{blue}}) / 32$$

Note: In the current version, the images to be registered have to have the same resolution and physical pixel/voxel size. An image is automatically re-

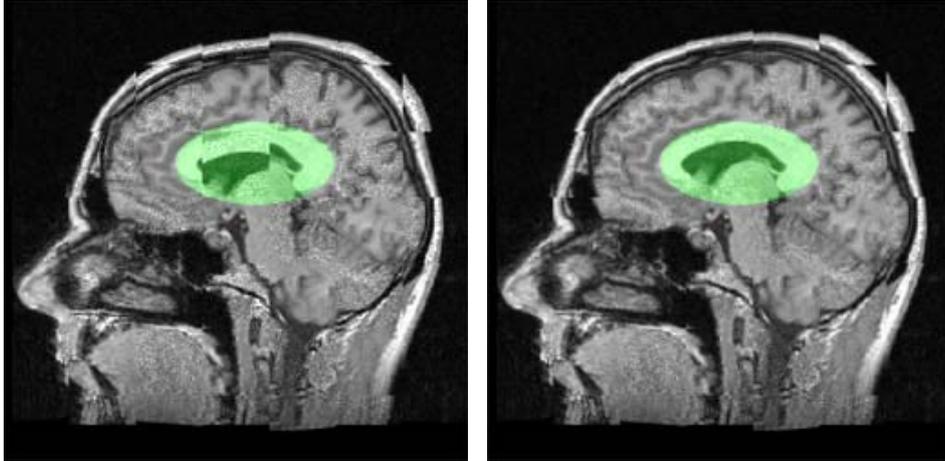


Figure 2: Registration with data mask.

sampled, if another image with different size is already loaded. Before loading a new set of images, the loaded images should be closed using *Close All Images* from the *File* menu. Please note, that the registration is considering the physical spacing and all parameters are given in millimeters. It is very important, to have the correct spacings in the header files of the images, especially in 3D with anisotropic voxels.

2.2 Data Masks

Optionally, a data mask can be loaded using *Load Data Mask* from the *File* menu. The data mask is a binary image of the same size as the images to be registered. Mask pixels/voxels with intensity zero are not considered when computing the similarity measure. The data mask can be removed using *Close Data Mask*.

2.3 Saving Results

The registration results can be saved in the MHD format only. The registered (deformed) image is saved using *Save Deformed Image* from the *File* menu. Additionally, the deformation field can be saved using *Save Deformation Field*. For the field, one MHD image is created per displacement dimension. An approximation of the inverse field can be saved using *Save Inverse Deformation Field*. The inverse field is computed from the original field using a thin-plate spline with equally distributed control points over the whole domain. The original field is then used as point correspondences in the evaluation of the thin-plate spline warping.

Note: The registration is using backward-warping for the image transformation. Thus, the displacements in the deformation fields will appear as pointing in the opposite direction of the actual motions.

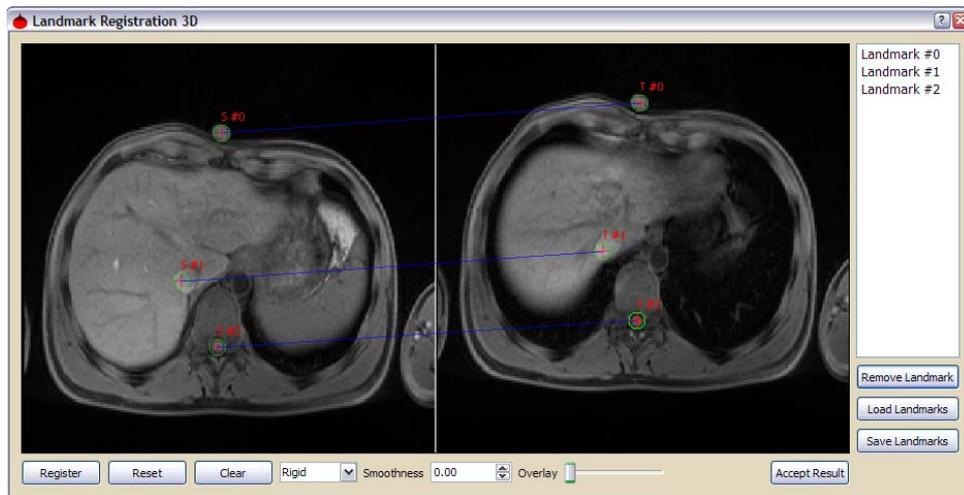


Figure 3: Landmark-based registration Tool.

2.4 Applying the Deformation Field

After registration, the deformation field can be applied to an external image using *Transform External Image* from the *Tools* menu. For instance, this can be useful if a segmentation saved in another file should be warped accordingly to the registration.

2.5 Landmark-based Pre-Registration

A simple landmark-based registration tool is provided. After loading two images, the landmark-based registration can be started by choosing Landmark Registration from the *Tools* menu. In 2D, at least four point correspondences have to be selected while eight points are needed for 3D. After choosing appropriate correspondences, the images can be pre-registered by an rigid, similarity, affine, or thin-plate spline warping. To accept the pre-registration, one has to click on *Accept Result* and the image is passed to the deformable registration application. It is possible to load and save correspondences from/into a simple text file. After accepting the result of the landmark registration, the corresponding displacement field, the inverse field, and the warped image are saved to the hard disk.

3 Registration Parameters

The registration tab is used to set the parameters of the registration algorithm. Each of the parameters will be explained in the following sections.

3.1 Control Grid Settings

The control grid defines the transformation model used for the registration. The size of the grid as well as the number of pyramid levels and the interpolation technique are application specific settings. In medical data, assuming

a global pre-registration, common parameters are for instance to start with a 20 mm control point spacing and a three level grid pyramid. When the algorithm switches to higher pyramid level, the control grid resolution as well as the image resolution is increased.

Set Distance (mm): Sets the initial control point distance in millimeters.

Image Levels: Sets the total number of image pyramid levels.

Grid Levels: Sets the number of control grid levels.

Min. Dimension: Sets the minimum resolution which will be preserved when generating an image pyramid.

Interpolation: Sets the interpolation method used to compute the dense deformation field.

Image Margin: Sets the boundary image margin in millimeters which is excluded from the similarity computation.

Update Mode: Selects the displacement field update mode.

Note: The combination of *compositional* updates and the *Link Maximum To Grid* option from the label set settings provides diffeomorphic deformation fields [7, 10]. The hierarchical registration is now controlled with two parameters – the image levels and the grid levels. The grid levels define the number of hierarchical iterations while the image levels only defines the number of levels used for the image pyramids. For example, if the number of image levels is set to 3 and the grid levels is set to 2, then the registration is performed only on the first two image pyramid levels with one subdivision of the grid. If the number of grid levels is set to 3 and number of image levels is set to 1, then the registration on all three grid levels is performed on the original resolution of the images. The minimum resolution parameter has an influence on the image pyramid. It is important to constrain the pyramid in certain cases. For example, oftentimes volumes have a very different resolution in z-direction. An $512 \times 512 \times 32$ image should not be rescaled in all dimension by a factor of 2. In fact, the small number of 32 slices should be preserved which can be enforced by setting the *Min. Dimension* parameter.

There are now three ways how the displacement field is updated after each iteration. *Additional* mode simply adds the dense displacement field update after each iteration. *Compositional* mode composes the updates by using linear interpolation. Both modes keep only the dense field in memory and the control points are reset to their initial position after each iteration. The *Consecutive* mode keeps the control grid deformation and updates are directly added to the control point displacements.

3.2 Label Set Settings

The label set definition is crucial for a successful registration. The labels define the search space for the displacements, but also influence the run time. The number of labels has to be chosen in a way to find a good compromise between speed and registration accuracy. In practice, it turned out that only a few labels are needed to perform a good registration. Due to the hierarchical approach in a coarse-to-fine manner, a sparse sampling of the search space is often sufficient.

Sparse/Dense Sampling: Sets the sampling method for the displacement space. A sparse sampling is generating displacements along the 8 main directions in 2D (vertical, horizontal, diagonal) and the 6 main directions in 3D (x-, y-, z-axis). The dense sampling is producing displacements equally distributed over an rectangular (2D) or cubic (3D) domain.

Link Maximum To Grid: Enables the displacement capture range to be stucked to the control grid resolution. The initial maximum displacement is set automatically to 0.4 times the control point spacing.

Maximum (mm): Sets the initial maximum value for the capture range manually. This is disabled if the link to maximum es enabled.

Steps: Sets the number of sampling steps. If sparse sampling is enabled, the total number of labels/displacements is $|L| = 8N + 1$ in 2D and $|L| = 6N + 1$ in 3D. If dense sampling is enabled, it becomes $|L| = (N + 1)^2$ respectively $|L| = (N + 1)^3$. The +1 accounts for the zero-displacement.

Label Factor: Sets the scaling factor for the capture range used after each optimization cycle. Thus, on each level the search space for the displacements is refined after each cycle which enables high accuracy registration results.

Note: The number of labels significantly influences the run time of the registration, since for each additional label the similarity measure has to be computed using the approximation scheme described in [4]. In practice, a sparse sampling is mostly sufficient when using the hierarchical approach and a sufficient number of iterations.

3.3 Optimization Settings

In the optimization settings, the actual cost function is defined which is subject to minimization.

Optimizer: Sets the discrete optimization method used for minimizing the cost function. Currently, only the FastPD [6] algorithm is available.

Data Cost: Sets the similarity measure. The following measures are currently available:

1. SAD: Sum of Absolute Differences

$$E(\mathcal{T}) = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} |I_1(\mathbf{x}) - I_2(\mathcal{T}(\mathbf{x}))| \quad (1)$$

where I_1 and I_2 are the target and source images and \mathcal{T} is the transformation.

2. SADG: Sum of Absolute Differences plus Sum of Gradient Inner Products

$$E(\mathcal{T}) = \frac{1}{|\Omega|} (1 - \gamma) \cdot \sum_{\mathbf{x} \in \Omega} |I_1(\mathbf{x}) - I_2(\mathcal{T}(\mathbf{x}))| + \gamma \sum_{\mathbf{x} \in \Omega} \left| \frac{\nabla I_1(\mathbf{x})}{|\nabla I_1(\mathbf{x})|} \cdot \frac{\nabla I_2(\mathcal{T}(\mathbf{x}))}{|\nabla I_2(\mathcal{T}(\mathbf{x}))|} \right| \quad (2)$$

where γ controls the balance between the two terms.

3. SSD: Sum of Squared Differences

$$E(\mathcal{T}) = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} (I_1(\mathbf{x}) - I_2(\mathcal{T}(\mathbf{x})))^2 \quad (3)$$

4. NCC: Normalized Correlation Coefficient

$$E(\mathcal{T}) = - \left| \frac{\sum_{\mathbf{x} \in \Omega} (I_1(\mathbf{x}) - \mu_1) (I_2(\mathcal{T}(\mathbf{x})) - \mu_2)}{\sqrt{\sum_{\mathbf{x} \in \Omega} (I_1(\mathbf{x}) - \mu_1)^2 \sum_{\mathbf{x} \in \Omega} (I_2(\mathcal{T}(\mathbf{x})) - \mu_2)^2}} \right| \quad (4)$$

5. NMI: A normalized version of mutual information (Entropy Correlation Coefficient) [8]

$$E(\mathcal{T}) = - \frac{2 \text{MI}(I_1, I_2 \circ \mathcal{T})}{\text{H}(I_1) + \text{H}(I_2 \circ \mathcal{T})} \quad (5)$$

where MI is the mutual information and H is the entropy. Uses the parameter *Histogram Bins* for defining the number of bins used for the estimation of the probabilities.

6. CR: Correlation Ratio [9]

$$E(\mathcal{T}) = -\text{CR}(I_1|I_2 \circ \mathcal{T}) \quad (6)$$

Uses the parameter *Histogram Bins* for defining the number of bins used for the estimation of the probabilities.

7. GRAD: Sum of Gradient Inner Products

$$E(\mathcal{T}) = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \left| \frac{\nabla I_1(\mathbf{x})}{|\nabla I_1(\mathbf{x})|} \cdot \frac{\nabla I_2(\mathcal{T}(\mathbf{x}))}{|\nabla I_2(\mathcal{T}(\mathbf{x}))|} \right| \quad (7)$$

8. CCGIP: Normalized Correlation Coefficient plus Sum of Gradient Inner Products

$$E(\mathcal{T}) = (1 - \gamma) \left(- \left| \frac{\sum_{\mathbf{x} \in \Omega} (I_1(\mathbf{x}) - \mu_1) (I_2(\mathcal{T}(\mathbf{x})) - \mu_2)}{\sqrt{\sum_{\mathbf{x} \in \Omega} (I_1(\mathbf{x}) - \mu_1)^2 \sum_{\mathbf{x} \in \Omega} (I_2(\mathcal{T}(\mathbf{x})) - \mu_2)^2}} \right| \right) + \gamma \sum_{\mathbf{x} \in \Omega} \left| \frac{\nabla I_1(\mathbf{x})}{|\nabla I_1(\mathbf{x})|} \cdot \frac{\nabla I_2(\mathcal{T}(\mathbf{x}))}{|\nabla I_2(\mathcal{T}(\mathbf{x}))|} \right| \quad (8)$$

where γ controls the balance between the two terms.

9. HD: Hellinger Distance

$$E(\mathcal{T}) = \sum_{\mathbf{x} \in \Omega} \left(\sqrt{p_1(I_1(\mathbf{x}))} - \sqrt{p_2(I_2(\mathcal{T}(\mathbf{x})))} \right)^2 \quad (9)$$

Uses the parameter *Histogram Bins* for defining the number of bins used for the estimation of the probabilities.

10. JRD: Jensen-Renyi Divergence [5] where γ controls the parameter α described in the paper :-). Uses the parameter *Histogram Bins* for defining the number of bins used for the estimation of the probabilities.

11. MI: Mutual Information [8]

$$E(\mathcal{T}) = -(\mathbb{H}(I_1) + \mathbb{H}(I_2 \circ \mathcal{T}) - \mathbb{J}\mathbb{E}(I_1, I_2 \circ \mathcal{T})) \quad (10)$$

Uses the parameter *Histogram Bins* for defining the number of bins used for the estimation of the probabilities.

12. JE: Joint Entropy

$$E(\mathcal{T}) = - \sum_{\mathbf{x} \in \Omega} p(I_1(\mathbf{x}), I_2(\mathcal{T}(\mathbf{x}))) \cdot \log(p(I_1(\mathbf{x}), I_2(\mathcal{T}(\mathbf{x})))) \quad (11)$$

Uses the parameter *Histogram Bins* for defining the number of bins used for the estimation of the probabilities.

Distance: Sets the penalty function used for regularization of the deformation field.

1. Pott's: A simple Pott's model penalizing deviations in the displacement assignments.
2. Truncated Absolute Difference: Penalizes the magnitude of vector differences between the displacements of two neighboring control points.
3. Truncated Quadratic Difference: Penalizes the squared magnitude of vector differences between the displacements of two neighboring control points.
4. Approximated Curvature Penalty: Penalizes the approximated second-order derivatives of the displacement field [2].
5. Distance Preserving: Penalizes changes in the distances between neighboring control points.

Projection: Sets the basis functions used for the back-projection of similarity measures (compare $\hat{\eta}$ in [4]). This parameter also selects the size of the image patch contributing to the similarity evaluation for each control point. If the projection is set to *Nearest*, only the closest pixels/voxels have an influence on the cost for moving a control point. A projection set to *Linear* considers all image points in the neighboring control grid cells, while *Cubic B-Splines* have the largest influence area (in 2D the 16 closest grid cells). The actual image point weightings defined by the selected basis function are only considered for point-wise similarity measures such as SSD and SAD. For the others, the projection parameter is only defining the patch size used for evaluating the local similarity (cp. [4]).

Incremental Regularization: Enables incremental regularization, meaning only the updates of the displacement field are penalized using the selected penalty function. If disabled, the whole displacement field over time is regularized.

Iterations: Sets the number of optimization cycles per pyramid level. After each of these cycles the label set is refined using the *Label Factor* parameter.

Gamma: Sets the γ parameter which is used for similarity measures such as SADG, CCGIP, JRD.

Lambda: Sets the λ parameter which controls the influence of the regularization term. This parameter is quite sensitive w.r.t. to the selected similarity measure, but also depends on the specific regularization method. A good strategy for finding a suitable value for λ is to start with a high value (e.g. 1000). If the result of the registration is only a rigid translation of the control grid, the value of λ can be successively halved until the control grid has enough flexibility for reasonable deformations.

Truncation: Sets the truncation threshold for the regularization terms. If set to zero, no truncation is used. In practice, a truncation of the regularization allows certain discontinuities in the regularization (cp. [1]).

Histogram Bins: Sets the number of bins used for the histograms. This parameter is needed for statistical similarity measures such as MI, NMI, CR, HD, JRD, JE.

3.4 Parameter File

The registration parameters can be loaded from and saved into simple text files. An example file with brief descriptions for the parameters can be found in the applications directory. The buttons *Load*, *Save*, and *Reload* above the label set settings can be used to load, save, and reload a parameter file.

References

- [1] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(11):1222–1239, Nov. 2001.
- [2] B. Glocker, N. Komodakis, N. Paragios, and N. Navab. Approximated curvature penalty in non-rigid registration using pairwise mrfs. In *5th International Symposium on Visual Computing (ISVC)*, Las Vegas, Nevada, USA, November 2009.
- [3] B. Glocker, N. Komodakis, N. Paragios, G. Tziritas, and N. Navab. Inter and intra-modal deformable registration: Continuous deformations meet efficient optimal linear programming. In *Information Processing in Medical Imaging*, Kerkrade, Netherlands, July 2007.
- [4] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios. Dense image registration through mrfs and efficient linear programming. *Medical Image Analysis*, 12(6):731–741, 2008.
- [5] Y. He, Y. He, A. Hamza, and H. Krim. A generalized divergence measure for robust image registration. *IEEE Transactions on Signal Processing*, 51(5):1211–1220, 2003.
- [6] N. Komodakis, G. Tziritas, and N. Paragios. Fast, approximately optimal solutions for single and dynamic mrfs. In *CVPR*, 2007.
- [7] S. Lee, G. Wolberg, K.-Y. Chwa, and S. Y. Shin. Image metamorphosis with scattered feature constraints. *Visualization and Computer Graphics, IEEE Transactions on*, 2(4):337–354, Dec. 1996.
- [8] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *Medical Imaging, IEEE Transactions on*, 16(2):187–198, April 1997.
- [9] A. Roche, G. Malandain, X. Pennec, and N. Ayache. The correlation ratio as a new similarity measure for multimodal image registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1115–1124. Springer-Verlag, 1998.
- [10] D. Rueckert, P. Aljabar, R. A. Heckemann, J. Hajnal, and A. Hammers. Diffeomorphic registration using b-splines. In *Proceedings of 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI 2006)*, October 2006.