VoxAlign Deformation Engine®

J.W. Piper, J.H. Richmond, A.S. Nelson, MD
MIM Software Inc.

Deformable Algorithms

The VoxAlign Deformation Engine provides a suite of deformable registration algorithms designed to handle the many scenarios that may be encountered clinically from same patient/same modality deformations to multi-modality deformations such as MR-CT or CT-CBCT.

Each DIR algorithm is initialized using a rigid registration that can be defined automatically or manually. Manual adjustment to the rigid registration allows the user to closely align a specific region of interest. Regions that are most closely aligned require less "searching" for corresponding anatomy by the algorithm and will usually result in the most accurate deformation.

Below is a description of each algorithm in the VoxAlign Deformation Engine followed by a table providing the specific information about each algorithm as recommended by AAPM TG-132.

Intensity-Based Deformable Registration

This is the default algorithm for same-modality deformable registration. It is a constrained, intensity-based, free-form deformable registration, which seeks to minimize intensity differences between the two images.

A coarse-to-fine, multi-resolution approach is used to define a grid of control points on the static image which are used to search for the best corresponding location in the target volume. Gross differences from the rigid registration are accounted for first and then refined to account for smaller local changes. The multi-resolution approach allows for good alignment even with large anatomical changes such as different shoulder or neck position and different phases of respiration. The final resolution of the computed deformation is no less than 3mm in each dimension.

The image matching metric minimizes intensity differences between the two images. The optimization strategy used is a custom modified gradient descent.

The deformation is regularized to avoid tears and folds in the deformation field. The regularization minimizes the effects of noise and incorrect correspondence while still providing a large degree of freedom for each control point to properly match the target volume.

Additional constraints are placed on the algorithm to further guide it toward reasonable results when executed on the same patient.

Normalized Intensity-Based Deformable Registration

This is the same algorithm as above, but with an intensity normalization scheme incorporated. A bi-linear model is fit to best normalize the intensities of one image to match those of the other. This is theoretically equivalent to having an unsupervised HU-ED calibration table for each CT and performing deformations in ED-scale instead of HU-scale.

Contour-Based Deformable Registration

There are two options for running a contour-based deformable registration: contour-only and hybrid (see below). The default option for contour-based deformable registration between non-CT modalities in MIM versions 6.8 and later is contour-only deformable registration. In contour-only deformable registration, MIM iteratively minimizes the difference in signed distance from the surfaces of contour pairs between the two images.

Hybrid Deformable Registration

This is the default option for contour-based deformable registration between CT scans in MIM versions 6.8 and later. Hybrid deformable registration is a combination of both contour-based and intensity-based DIR described above. The registration metric minimizes both the intensity differences between the two images and the differences between corresponding contour surfaces defined on the two images. The influence of the contour surface matching is inversely proportional to the distance from the surface.

Multi-Modality Deformable Registration

This was developed mainly for direct MR-to-CT deformation. However, it is a general use multi-modality free-form deformation that uses a feature similarity scoring metric. It aims to maximize the correspondence of high-dimensional feature descriptors computed by evaluating each image voxel in the context of its neighboring voxels.

Convert Local Alignments into Deformable Registration

This option considers all of the locked Reg Refine™ local alignments and combines them into an overall deformable registration. It does not look at the intensities of the two scans. The combination uses a Gaussian mixing model to spatially weight the contributions of each local rigid alignment. Notice that any part of the secondary image that is too distant from a locked alignment will be masked from the resulting deformed image. The reason for this is to ensure sufficient local rigid alignments are made.

Reg Refine

Reg Refine is an input into multiple DIR algorithms. It allows the user to define local rigid alignments to provide additional information to help guide the deformation algorithm near these areas. It can be used iteratively to execute a DIR, evaluate the local DIR accuracy, and suggest local alignments to improve the DIR result until an optimal alignment is achieved.

	Intensity-Based	Normalized Intensity	Convert-Local	Contour-Based	Hybrid-Based	Multi-Modality
Similarity Metric	Sum of squared differences.	Sum of squared differences (after intensity normalization using a bi-linear model).	None	Sum of squared differences between signed distance functions.	Combination of (a) sum of squared differences between signed distance functions representing contours defined on the two images and (b) intensity differences between the two images.	High-dimensional feature descriptor based on differences from neighboring voxels.
Transformation & Regularization Model	Free-Form	Free-Form	Gaussian Mixture Model	Free-Form	Free-Form	Free-Form
Optimization Method	Modified Gradient Descent	Modified Gradient Descent	Analytical Solution	Modified Gradient Descent	Modified Gradient Descent	Gauss-Newton method
Knobs	Smoothness Dynamic regularization Reg Refine	Smoothness Dynamic regularization Reg Refine	Reg Refine		Smoothness Dynamic regularization	Smoothness Reg Refine
Notes	Constrained, intensity-based, free-form deformable registration algorithm.	Constrained, intensity-based, free-form deformable registration algorithm.		Minimizes surface differences.	Minimizes both sum of squared differences and surface differences.	
Example Use Cases	Same modality: CT-CT, CBCT-CBCT	Same modality: MR-MR, CT-CBCT	Multi-modality cases where other DIR algorithms fail.	Any 3D modalities	CT-CT, CT-CBCT, CBCT-CBCT	Multi-modality, MR-CT, different sequence MR, CT-CBCT, Contrast-Non-Contrast CT

DIR Profiles

A single algorithm is unable to perform optimally for all of the clinical scenarios in which deformations are needed. In MIM version 6.6.0, DIR Profiles were introduced to provide the ability to define the optimal parameters to use for different clinical scenarios.

DIR Profiles provide the ability to select the "Registration Type" as well as the level of smoothing ("Smoothness Factor"), choose to use intensity "Normalization," and choose to use "Dynamic regularization."

Registration Types

These registration defaults were developed by re-optimizing the intensity-based DIR method using real clinical CT and CBCT data for a variety of treatment areas and representing a variety of use cases. The new intensity-based Registration Types are: Same Subject (MIM 6.6.0+) and CBCT to CT (MIM 6.6.0+). The previous intensity-based parameters are also provided: Same Subject (MIM < 6.6.0) and CBCT (MIM < 6.6.0). The new multi-modality algorithm has also been included as a Registration Type in the DIR Profiles.

Smoothness Factor

The smoothness factor provides the ability to control the smoothness of the deformation. Smaller numbers (less smoothing) allow the algorithm to have more freedom to move to account for very local changes in anatomy. However, if there is insufficient information to guide the algorithm, this can lead to the result being driven by noise or other artifactual information in the image. Conversely, as the smoothing is increased (higher number), this causes the registration to have less freedom to account for very local changes. However, for use cases where there is insufficient information to guide the deformation, smoothing will help to prevent unrealistic deformations from occurring in these regions of little detail.

Normalization

A bi-linear model is fit to normalize the intensities of one image to match those of the other. This is used for cases where intensities between images are inconsistent (e.g., CT-CBCT, MR-MR, etc.).

Dynamic Regularization

This is used to provide less smoothing in areas with sufficient information to help drive the deformation and more smoothing in areas with less information to prevent unrealistic deformations from occurring.

Reg Reveal® and Reg Refine

The most important question in DIR is how the deformation worked on an individual patient. Every patient is different, and evaluating the accuracy of DIR can be critical when performing dose accumulation or DIR for target contouring using diagnostic images.

Reg Reveal was developed for this purpose and is currently the only tool available for the specific purpose of efficient quality assurance (QA) of DIR. Reg Reveal allows the user to interrogate the registration in specific regions of interest and draw conclusions about its accuracy. Such a deformable registration evaluation tool, which allows the user to answer the essential question of whether a DIR algorithm correctly identifies corresponding anatomy for a given patient, should be considered an essential QA tool when implementing DIR in the clinic. Reg Refine allows the user to influence the registration algorithm to achieve a more accurate result.