



Hubless keypoint-based 3D deformable groupwise registration

Rémi Agier, Sébastien Valette, Razmig Kéchichian, Laurent Fanton, Rémy Prost

► To cite this version:

Rémi Agier, Sébastien Valette, Razmig Kéchichian, Laurent Fanton, Rémy Prost. Hubless keypoint-based 3D deformable groupwise registration. Medical Image Analysis, 2019, 10.1016/j.media.2019.101564 . hal-01977879

HAL Id: hal-01977879

<https://hal.science/hal-01977879>

Submitted on 20 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Hubless keypoint-based 3D deformable groupwise registration

R. Agier^a, S. Valette^{a,*}, R. Kéchichian^a, L. Fanton^{a,b}, R. Prost^a

^a *Université de Lyon, INSA Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne
CNRS, Inserm, CREATIS UMR 5220, U1206, F69621 LYON, France*

^b *Hospices Civils de Lyon, GHC, Hôpital Edouard-Herriot, Service de médecine légale,
69003 LYON, FRANCE.*

Abstract

We present a novel algorithm for Fast Registration Of image Groups (FROG), applied to large 3D image groups. Our approach extracts 3D SURF keypoints from images, computes matched pairs of keypoints and registers the group by minimizing pair distances in a hubless way i.e. without computing any central mean image. Using keypoints significantly reduces the problem complexity compared to voxel-based approaches, and enables us to provide an in-core global optimization, similar to the Bundle Adjustment for 3D reconstruction. As we aim to register images of different patients, the matching step yields many outliers. Then we propose a new EM-weighting algorithm which efficiently discards outliers. Global optimization is carried out with a fast gradient descent algorithm. This allows our approach to robustly register large datasets. The result is a set of diffeomorphic half transforms which link the volumes together and can be subsequently exploited for computational anatomy and landmark detection. We show experimental results on whole-body CT scans, with groups of up to 103 volumes. On a benchmark based on anatomical landmarks, our algorithm compares favorably with the star-groupwise voxel-based ANTs and NiftyReg approaches while being much faster. We also discuss the limitations of our approach for lower resolution images such as brain MRI.

Keywords: Groupwise registration, keypoints

1. Introduction

Registration of several images together, also known as groupwise registration, is nowadays most often carried out for human brain studies (Jenkinson et al., 2002). Whole body studies remain rare as they still raise significant problems (Xu et al., 2016). Advances in imaging techniques are constantly increasing the number and size of 3D images in hospital databases, hence the need for

*Corresponding author

Email address: sebastien.valette@creatis.insa-lyon.fr (S. Valette)

low complexity groupwise image registration techniques. Instead of dense voxel based registration, sparse keypoint matching, extracted from the images is a promising approach. Keypoints are extracted with their location and description vector. Points from different images are paired by comparing their description vectors. This is a great challenge when shifting from inpatient registration to interpatient registration, as the human anatomy exhibits a large variability.

In this paper we propose a novel groupwise registration approach aimed at large image databases which is able to register high resolution images, such as whole body CT scans shown in Figure 9. As human anatomy exhibits high variability (e.g. men vs. women), our algorithm is hubless, i.e. it does not use any central reference during registration. Inspired by advances in the computer vision field, our algorithm exploits keypoint detection and matching, which bring speed and robustness. Some paired points are erroneous (outliers) and we devised an algorithm for robustness against these outliers, even when they are very frequent. Potential applications range from computational anatomy to forensic anthropology and robust patient-specific model construction. An example of forensic application is probabilistic sex diagnosis using worldwide variability in hip-bone measurements (Murail et al., 2008). Moreover, the ever-increasing amount of medical images stored in Picture Archiving and Communication Systems (PACS) in hospitals offers a great opportunity for big data analysis. Screening huge image groups could improve diagnosis and could be of critical help for computational anatomy. Recent advances in machine learning (Michalski et al., 2013) or atlas based-approaches (Iglesias and Sabuncu, 2015) have pushed for efficient big data analysis tools. Yet the modest size of current annotated medical image databases limits the use of these approaches for medical imaging. As a consequence, whole-body groupwise registration could bridge the gap between big data and organ localization, segmentation, and computational anatomy. One known limitation for our approach is that the number of extracted keypoints should be as high as possible to obtain the best accuracy.

The paper is organized as follows: section 2 presents previous works related to registration, keypoints and Bundle Adjustment. Section 3 summarizes the key contributions of our hubless approach. Section 4 explains our approach in detail while section 5 shows results obtained with our approach, as well as comparisons with both NiftyReg and ANTs star-groupwise algorithms. We conclude the paper in section 6.

2. Related works

2.1. Image registration

Image registration, also known as image matching, fusion or warping, consists of finding a transform τ between two or more images, mapping any point p from a source image to a position $\tau(p)$ in the target image. It is a crucial step in many medical applications, such as longitudinal studies (Scahill et al., 2003), radiotherapy planning (Keall et al., 2005), brain studies (Jenkinson et al., 2002), atlas-based segmentation (Lötj  nen et al., 2010), image reconstruction (Huang

et al., 2008) and microscopy (Vercauteren et al., 2007). We present here a brief state of the art, and we refer the reader to Brown (1992); Maintz and Viergever (1998); Sotiras et al. (2013) for thorough reviews of existing approaches.

2.1.1. Pairwise & Groupwise registration

Pairwise registration consists of matching one moving image to a fixed image. Groupwise registration consists of registering a whole set of images together.

One can categorize groupwise registration approaches using graph theory, as shown in Figure 1, where vertices represent images and edges represent transforms. Then, transforming any image to any other image requires at least a graph spanning all vertices. Then at least $n - 1$ transforms are needed for n images.

- Star groupwise methods register each image of the set against a reference image. In this case the graph is a star graph (Figure 1(a)). The output is a set of half transforms τ_i , mapping each image i to the common space. After registration, the transform $\tau_{i \rightarrow j}$ from one image i to another image j is a composition of the two half transforms τ_i and τ_j : $\tau_{i \rightarrow j} = \tau_j^{-1} \circ \tau_i$. These methods are biased by the choice of the initial reference or template image. This bias takes two forms : the intensity bias and the shape bias. Guimond et al. (2000) reduced these biases by iteratively updating the reference image, taking into account the average of each transformed image, and the average of image transformations. Later, Joshi et al. (2004) improved the theoretical foundations of this approach via diffeomorphism. An example of application for brain pediatric studies is shown in Fonov et al. (2011). The publicly available softwares ANTs(Avants et al., 2008), Elastix(Klein et al., 2010) and NiftyReg(Modatt et al., 2008) belong to the star groupwise category.
- Tree groupwise methods use a spanning tree graph and multiple references (Wu et al., 2011), as shown by Figure 1(b). Variability can then be distributed across several references instead of only one.
- Hubless methods use an abstract common space, but do not use any reference image: while star and spanning tree graphs are minimal graphs, hubless methods use a dense graph to benefit from more constraints, as shown by Figure 1(c). When registering n images, up to $n(n - 1)/2$ observations can be used. Observations can be pairwise registrations, or local matches. Hamm et al. (2010) proposed GRAM, a hubless registration method, based on manifold learning. This algorithm requires dense pairwise registration of all image pairs, which is very time consuming. As an example, the authors report a processing time of 24 hours on a cluster to register 416 low resolution images ($68 \times 56 \times 72$). Ying et al. (2014) proposed a similar approach, using geodesic graph shrinkage. A recent work on groupwise nonrigid registration was also proposed by Wu et al. (2012), but is dedicated to the processing of brains acquired with MRI, as it requires a segmentation of the image into White Matter, Grey Matter




Figure 1: Different groupwise registration methods. Images are outlined in red, reference images are outlined in dotted black, transforms in dashed orange and observations in black. For dense registrations, observations are generally voxelwise difference, cross-correlation or mutual information. For sparse registrations, observations are usually keypoint match distances. Transforms are optimized via the minimization of an energy function driven by the observations. (a) star groupwise registration (1 reference, n transforms and n observations). (b) tree-groupwise registration (multiple references, $n-1$ transforms, $n-1$ observations). (c) hubless registration, with no reference image but one abstract common space, n half-transforms and at most $\frac{n(n-1)}{2}$ observations.

and Ventricular Cerebrospinal Fluid. More recently, the approach proposed in Agier et al. (2016) uses no reference at all. This allows handling partially matching body parts and high variability but is restricted to rigid transforms.

Groupwise registration has most often been used for brain analysis, as the brain shape exhibits low variations between subjects overall. Whole body studies such as carried out by Suh et al. (2011) remain rare. The important advantage of hubless approaches becomes clear when averaging a group of images is a problem. A good example is to register full body images, where computing the average between men and women is inconsistent in several regions.

2.1.2. Rigid & Deformable Registration

Rigid registration (Ashburner et al., 2003) consists of computing a linear transform between two images. The transform can be a translation or a more expressive one such as a combination of translation, rotation, or scaling. Rigid transforms can be used for Procrustes Analysis (Schönemann, 1966), Generalized Procrustes Analysis (Bartoli et al., 2013), mosaic assembly or as an initialization step for nonrigid transform. Nonrigid (deformable) registration consists of finding a free-form transform able to express local variability, which cannot be done with a rigid transform. Nevertheless, the transform cannot be arbitrary, and should generally be smooth and invertible.

2.1.3. Dense & Sparse Registration

Registration approaches can also be split in two classes, depending on how they process the input data:

Dense registration estimates displacement on the entire image domain (Avants et al., 2008; Modat et al., 2008, 2014). The most common approaches use intensity difference between images but there are many other criteria such as cross-correlation (Lewis, 1995) or mutual-information (Pluim et al., 2003). The output is a displacement vector for each input voxel. The key difficulties with dense registration are high computational cost and an ill-posed optimization problem, which generally needs explicit regularization (Robinson and Milanfar, 2004). As an example, the Advanced Normalization Tools (ANTs) framework (Avants et al., 2008) registers two thoracic images in about 2 hours on our test machine.

Sparse registration, instead of working on a dense grid, uses only point sets. Compared to dense registration, sparse registration is much faster (Allaire et al., 2008; Cheung and Hamarneh, 2009) but may yield less accurate results because points do not always span the whole space. For nonrigid registration, sparse transforms need an interpolant, such as splines (Szeliski and Coughlan, 1997) or radial basis functions (Fornet et al., 2001). The interpolant inherently provides a smooth solution which helps regularizing the solution. Point sets can be of different types:

- Reference landmarks which are distinguishable anatomic structures manually placed on the image by an expert. This task is usually carried out by physicians and can be time consuming. These are used in (Wang et al., 2006; Li et al., 2012)
- Reference landmarks automatically placed on the image by an algorithm, such as proposed by Zhu et al. (Zhu et al., 2013)
- Vertices of a surface mesh representing the object to register, such as proposed in (Rasoulzadeh et al., 2012)
- Keypoints, automatically detected and placed in the image. One important difference between reference landmarks and keypoints is that keypoints are not defined by anatomical definitions but by mathematical properties of the image. Hence, although it is interesting to find the same keypoints in each image, this is never the case in practice, and the number of extracted points from each image is not a priori known. Next section (2.2) gives more details on keypoints.

Some works use a combination of both dense and sparse approaches. As an example, Wu et al. (2012) uses driving voxels to speed up a dense method.

2.2. Keypoints

During the last decades keypoints (Harris and Stephens, 1988) have successfully been used for object recognition (Lowe, 1999), action recognition (Wang et al., 2011), robotic navigation with Simultaneous Localization and Mapping (Karlsson et al., 2005) and panorama creation (Anguelov et al., 2010). They aim at being fast while reducing the amount of data to process, mainly to deal with

real time processing or tasks involving large amounts of data. Their suitability to medical imaging has been evaluated in (López et al., 1999), and various applications have been proposed in this context such as image annotation (Datta et al., 2005), retrieval (Zheng et al., 2008) and segmentation (Kéchichian et al., 2014; Wachinger et al., 2018). Initially designed for 2D images, keypoint approaches have been extended to process 3D medical images by (Allaire et al., 2008; Agier et al., 2016; Cheung and Hamarneh, 2009; Rister et al., 2017).

The keypoint extraction pipeline is the following: first, for each image, a detector extracts locations exhibiting important features such as corners (first derivative analysis (Harris and Stephens, 1988)) or blobs (second derivative), most often via fast approximations of these derivatives. As an example, Scale Invariant Feature Transform (SIFT) uses difference of Gaussians (Lowe, 2004) and Speeded Up Robust Features (SURF) use integral images (Bay et al., 2006). Afterwards, the neighborhood of each location is summarized into a compact feature vector, called descriptor. SIFT populates the feature vector with orientation histograms, while SURF stores local texture information.

Efficient registration can be performed as follows: (1) keypoints are extracted from the input images, with locations and descriptions. (2) Keypoints from different images are paired by comparing their feature vectors. Note that the number of keypoints may vary between images. This results in a set \mathcal{M} of paired points. (3) Registration is performed by minimizing the distance between all paired points of this set (Mikolajczyk and Schmid, 2005).

The use of paired points combined with sparse registration approaches leads to very fast registration, in opposition to dense registration. Moreover, as keypoints have to be extracted only once per image, repeated registration with the same reference image is very efficient.

Finally, keypoints are rarely used for groupwise registration, most of the previous approaches use anatomical reference landmarks to perform groupwise registration. Few papers deal with keypoints, such as Zhang et al. (Zhang and Cootes, 2012). Note that (Zhang and Cootes, 2012) is restricted to 2D images.

2.3. Outlier rejection

A drawback of using keypoints is the significant number of mismatches (outliers) in the paired point set \mathcal{M} . Outliers are pairs of keypoints with similar descriptors describing regions which are actually different and therefore should not be paired. Decreasing the image quality (using low end cameras and webcams used in consumer electronics devices) potentially increases the outlier ratio to more than 30% of the set (Läbe and Förstner, 2004) (37% in (Mikolajczyk and Schmid, 2002)).

In Rangarajan et al. (1997), a spatial mapping and the one-to-one correspondences (or homologies) between point features extracted from the images try to reject non-homologies as outliers. In Chui and Rangarajan (2003), an alternating optimization algorithm successively updates the points correspondence matrix and the transformation function while gradually reducing the temperature in a deterministic annealing. The points correspondence matrix handles outliers.

In Myronenko and Song (2010), a probabilistic method called Coherent Point Drift (CPD) fits the first point set centroids as a mixture of Gaussian with an additional uniform distribution to account for outliers point set centroids to the data points by likelihood maximization. The method forces the centroid to move coherently as a group to preserve the topological structure.

M-estimators (Huber, 2011) provide a good way to mitigate the influence of outliers in this context. During optimization, M-estimators provide an adaptive weighting function which adjusts the contribution of each pair by computing robust statistics on \mathcal{M} . Each point pair can then be classified as inlier or outlier (soft classification) and weighted accordingly. For this aim, one needs an estimation of the inlier contribution variance, computed via Median Absolute Deviation (MAD). Note that the maximum outlier ratio that MAD can process (the breakdown point (Donoho and Huber, 1983)), is 50%. Hence the breakdown point of the M-estimator is also 50%.

2.4. Bundle Adjustment

Bundle adjustment (BA) is a class of computer vision approaches which allows reconstruction of a unique 3D scene from multiple 2D views by back projection of keypoints in the 3D scene and global optimization (Triggs et al., 2000). It allows efficient tracking even with low-end cameras (Karlsson et al., 2005). Recent BA implementations can manage large amounts of data and successfully reconstruct large parts of a city (Frahm et al., 2010).

BA jointly optimizes point positions, camera positions and camera parameters. Because of (1) the non-linearity of projective geometry (Hartley and Zisserman, 2003), (2) mismatches and (3) the M-estimator, optimization is high-dimensional and non-convex (Dauphin et al., 2014). As a consequence, particular attention has to be paid to the optimization method. Algorithms such as Levenberg-Marquart optimization (Moré, 1978) can be used to avoid converging to a local minimum. BA also needs a good initialization to remain close to the global minimum at any time.

3. Challenges and contributions

Our paper provides a unified solution to register in a nonrigid way a whole image bundle at once. It can be seen as a complete graph registration because each image is linked with all other images (figure 1(c)). Several challenges need to be addressed:

- *Amount of data:* As we want to process many images in reasonable time, we cannot use dense registration: each image from our test database contains about 512^3 voxels, each voxel being encoded with 32 bits. Depending on the experiments, 20 to more than 100 images are used, representing more than 50GB of image data. Current approaches process a lower number of images, as in (Shattuck et al., 2008). On top of a hubless approach

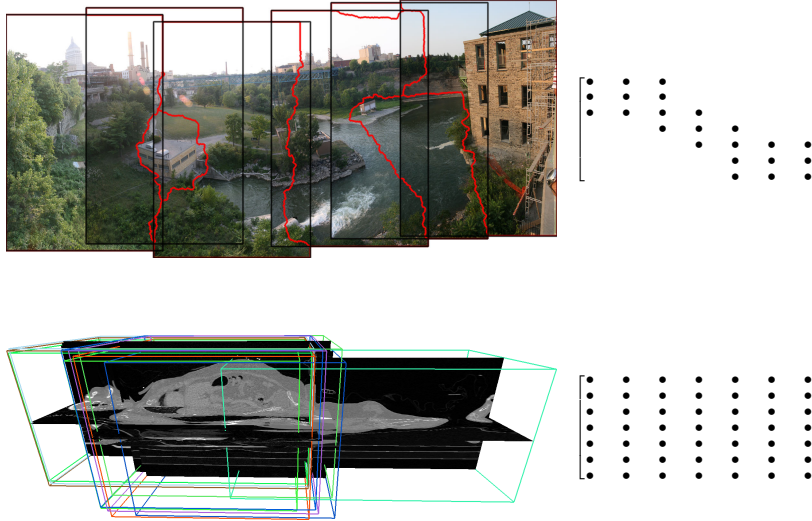


Figure 2: Sparsity induced by a panoramic reconstruction with BA (top) compared to our application (bottom). Overlap matrices of each problem are given on the right. The top image is provided by Noso1, via Wikimedia Commons.

as in (Wu et al., 2012), we need a compact solution which avoids reading voxel data and allows us to manage many images in an in-core way.

- *Repeatability & Outliers:* We have to overcome inter-patient variability. When using keypoints, the rate of outliers increases from 30% for similar patients up to 70% in case of different patients (see section 5.3.3). Two main reasons stand out, both induced by inter-patient variability : (1) match selectivity needs to be decreased to find more matches between differing patients, (2) match consistency cannot be enforced (if three points A, B, C are linked by two matches : AB and BC, A and C must match). As a result, the rate of outliers can exceed the M-estimator breakdown point (50%), hence the need for a more robust estimator.
- *Computational complexity:* The analogy with BA is not perfect. BA processes a set of 2D images distributed in a 3D space. As each image overlaps with few other images, BA can be significantly simplified to a band matrix problem. In our case, we need to register 3D images in a 3D space, and the overlap between images is much higher: the images are all located within the human body, but for different patients. This results in a dense matrix problem, illustrated by Figure 2.

To overcome these challenges, we introduce in this paper three main contributions:

- *Keypoint-based hubless registration:* As input data, we use only the keypoints extracted from the images. We still use an abstract common space




Figure 3: The three main steps in our algorithm: (1) Keypoints are extracted from each image. (2) Matching constructs a list of point pairs from the point sets. (3) Registration optimizes a transformation for each image in order to minimize point pair distances in the common space.

and half-transforms as in star groupwise registration, but we do not need any central reference data : our optimization is driven only by inter-image registration, as shown in Figure 1(c). Half-transforms are represented by spline pyramids, and we use the graph linking all images together (the complete graph) for the optimization. This compact framework can register 100 images with a memory footprint of only 10GB.

- *EM-weighting*: To manage a ratio of outliers that can exceed 50%, we devise an Expectation-Maximization algorithm which explicitly estimates inlier and outlier distributions. Moreover, experimental results show that the proposed EM-weighting yields better selectivity against outliers than the M-estimator. As a result, EM-weighting allows our algorithm to converge, while the M-estimator failed to provide convergence.
- *Efficient optimization*: We propose a fast energy function minimization algorithm suited to our nonrigid problem. With this algorithm, our 24-core testing computer registers 20 images (512x512x400 voxels each) in about 10mn on our test machine. A laptop needs about 30mn.

4. Method

Our algorithm proceeds in three steps : (1) 3D SURF keypoint extraction from all input images, (2) points pairing according to their descriptor, (3) half-transforms optimization by minimizing the distance between paired points, while aiming at discarding outlier influence. Figure 3 shows a block diagram of the three steps.

4.1. Half transforms driven by keypoints

We consider a set \mathcal{I} of n 3D images. In order to get rid of dense computation at the earliest point in the process, we extract for each image i the 3D SURF keypoint set $\mathcal{P}^i = \{p_a^i\}$ as done by Agier et al. (2016). For simplicity, we define $\mathcal{P} = \bigcup \mathcal{P}^i$. The entire algorithm uses physical coordinates (in millimeters),

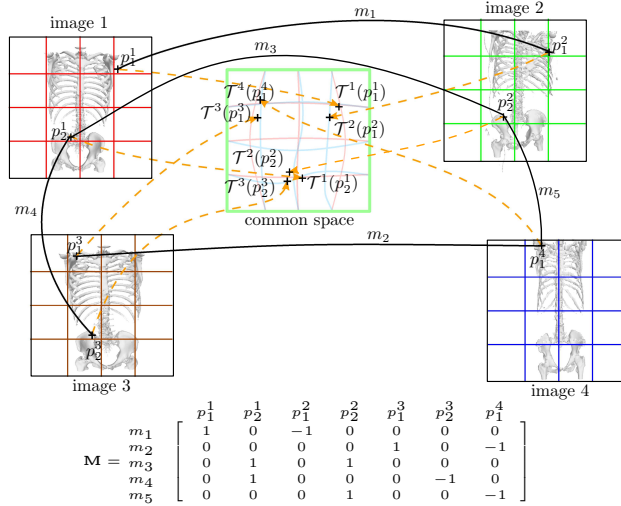


Figure 4: Example with 4 input images. For clarity purposes, only bone structures are displayed. Overprinted on each image i : the grid on which lie the splines used to represent the nonrigid transform τ^i . In the middle : the common space, which is used to measure match distances. Similarly to Figure 1(c), matches are drawn in solid black while point transforms are drawn in dashed orange. In this simple example, 7 keypoints were extracted : $\mathcal{P} = \{p_1^1, p_2^1, p_1^2, p_2^2, p_1^3, p_2^3, p_1^4\}$, and 5 matches were found: $\mathcal{M} = \{(p_1^1, p_1^1), (p_1^3, p_1^4), (p_2^1, p_2^2), (p_2^1, p_2^3), (p_2^2, p_1^4)\}$. The match (p_2^2, p_1^4) is an outlier. \mathbf{M} is the match matrix.

and is invariant to image sampling conditions. We compute the set of matched keypoints $\mathcal{M} = \{(p_a^i, p_b^j)\}$ for all image pairs (i, j) , by combining several criteria: feature vector distance, nearest neighbor ratio and Laplacian sign as in Bay et al. (2006) and scale difference. As explained earlier in the paper, \mathcal{M} contains outliers. Optimizing the set of half transforms $\mathcal{T} = \{\tau^i\}$ is done by minimizing the match (p_a^i, p_b^j) distance-errors :

$$d(p_a^i, p_b^j) = \|\tau^i(p_a^i) - \tau^j(p_b^j)\| \quad (1)$$

This L_2 distance is the only measure computed in the abstract common space. We use tensor products of Uniform Cubic B-splines to interpolate the half-transforms τ^i , driven by c control points placed on a rectilinear grid. Each half-transform τ^i is expressed as: $\tau^i : p \in \mathbb{R}^3 \mapsto p + \mathbf{b}(p)\mathbf{x}^i$, where the j^{th} element of $\mathbf{b}(p)$ is the evaluation of the j^{th} spline basis function at point p ; \mathbf{x}^i is the $c \times 3$ spline coefficient matrix which represents the control point displacements (one column for each coordinate). \mathbf{x}^i are unknowns we want to optimize for each image. Figure 4 gives an overview of our framework, a simple example with 4 input images, 7 keypoints and 5 matches.

4.2. Outliers - EM-weighting

An efficient way of mitigating the presence of outliers in the set of paired points is to use robust statistics such as M-estimators (Fox, 2002). The usual

assumption made with M-estimators is that the probability distribution of the distance (1) is a combination of one normal distribution (the inlier contribution), and one uniform distribution (the outlier contribution). M-estimators need to estimate inlier variance, the tuning parameter, which provides a way to balance the contribution of each distance $d(p_a^i, p_b^j)$ (1) to the optimization criterion. Smaller values of the tuning constant increase robustness to outliers.

In our case, the underlying distributions follow different laws:

- **Inliers** are norms of Gaussian random vectors of \mathbb{R}^3 . They follow a χ distribution with 3 degrees of freedom, also known as Maxwell distribution.
- **Outliers** are distances between random vectors of \mathbb{R}^3 following uniform distributions. We propose to approximate their distribution using also a Maxwell distribution.

Then, the probability density function of the distance can be expressed on each image i as:

$$P^i(d) = r^i f(d, s_1^i) + (1 - r^i) f(d, s_2^i) \quad (2)$$

$$f(d, s) = \sqrt{\frac{2}{\pi}} \frac{d^2}{s^3} e^{-\frac{d^2}{2s^2}}$$

where f is the Maxwell probability density function, s_1^i and s_2^i are respectively the inlier and outlier scale parameters (the standard deviation of all vector distance coordinates with zero mean) of the two Maxwell laws and r^i is their mixing ratio. Instead of estimating the inlier variance, we directly estimate for each image i the parameters $\theta^i = (s_1^i, s_2^i, r^i)$ of $P^i(d)$ using the Expectation-Maximization (EM) algorithm (Dempster et al., 1977). As we explicitly estimate the Maxwell laws, we can handle more than 50% of outliers as we do not rely on the median operator. Following Bayes' rule, we can infer for any match distance-error d the probability of belonging to the inlier class I :

$$P(I|d, \theta^i) = \frac{r^i f(d, s_1^i)}{r^i f(d, s_1^i) + (1 - r^i) f(d, s_2^i)} \quad (3)$$

This probability can be used as a weighting function to inhibit outlier contributions. But as each match (p_a^i, p_b^j) links two images, two probabilities can actually be computed for each match: $P(I|(p_a^i, p_b^j), \theta^i)$ and $P(I|(p_a^i, p_b^j), \theta^j)$, using respectively image i and image j as statistical context. As a symmetric criterion is preferable, we select the minimum value as a weight for each match:

$$w(p_a^i, p_b^j) = \min[P(I|d(p_a^i, p_b^j), \theta^i), P(I|d(p_a^i, p_b^j), \theta^j)] \quad (4)$$

The minimum is a good heuristic since it increases selectivity against outliers. It has been observed in our experiments that it is more effective to wrongly reject more inliers than wrongly accept more outliers. Figure 5 shows example data as well as a comparison between EM-weighting and M-estimators, showing that our approach is more robust and more selective than the M-estimator.




Figure 5: EM-weighting and M-estimators. Top : synthetic distance distribution, with inlier distribution in light blue and outlier distribution in orange. The dark-blue and red curves represent their approximations with Maxwell distributions. Middle : Real world example, with more than 50 % of outliers. The gray histogram is computed for one image of the VIS-CERAL data set after registration. The dark-blue and red curves respectively represent the estimation of inlier and outlier distributions with Maxwell distributions. Bottom : Comparison of weighting functions computed with the same real world data. Blue: incorrect standard deviation estimation leads the M-Estimator to include a lot of outliers. Green: M-estimator with corrected standard deviation. Red: our EM-weighting yields better selectivity.

4.3. Energy minimization

4.3.1. Initialization

Before non-rigid registration, we optimize a 6-parameter transformation (translation + anisotropic scale) λ^i defined for each image i as:

$$\lambda^i : p \mapsto \lambda^i(p) = s^i \circ p + t^i \quad (5)$$

Where s^i and t^i are vectors of \mathbb{R}^3 representing the anisotropic scale and translation, respectively. The \circ symbol stands for the component-wise Hadamar product. Optimization is carried out using an iterative mean and variance matching algorithm : for each iteration and each image i , the scale s_i is updated to \tilde{s}_i so that the component-wise variance v_a of keypoints p_a^i in i gets closer to the variance v_b of their paired keypoints :

$$\tilde{s}_i = ((v_{bx}/v_{ax})^{\frac{\gamma}{2}}, (v_{by}/v_{ay})^{\frac{\gamma}{2}}, (v_{bz}/v_{az})^{\frac{\gamma}{2}})^T \quad (6)$$

where γ is the update coefficient, lower than 1, and v^T is the transpose of vector v . The translation t_i is updated to \tilde{t}_i to displace in the common space the mean \bar{p}_a of the keypoints p_a^i closer to the mean \bar{p}_b of their paired keypoints:

$$\tilde{t}_i = t_i + \gamma(\bar{p}_b - \bar{p}_a) + \bar{p}_a \circ ((1, 1, 1)^T - \tilde{s}_i) \quad (7)$$

Note that in order to reduce the influence of outliers, paired points are weighted by our EM-weighting scheme $w(p_a^i, p_b^j)$ (4) during computation of the

variances (v_a and v_b) and the mean values ($\overline{p_a}$ and $\overline{p_b}$). In our experiments, we carried out 50 iterations with $\gamma = 0.5$, and update the weights $w(p_a^i, p_b^j)$ once every 10 iterations.

4.3.2. Deformable registration

Once linear registration is computed, we propose to find the best set of B-Spline transforms \mathcal{T}^* by minimizing for each keypoint the mean square distance to its matching points, each match contribution (1) being weighted according to EM-weighting (4):

$$\mathbf{E}(\mathcal{T}) = \sum_{i=1}^n \sum_{p_a^i \in \mathcal{P}^i} \left(\frac{1}{|\mathcal{N}(p_a^i)|} \sum_{p_b^j \in \mathcal{N}(p_a^i)} w(p_a^i, p_b^j) d(p_a^i, p_b^j)^2 \right) \quad (8)$$

where $|\mathcal{N}(p_a^i)|$ is the number of points matching with p_a^i .

However, $n-1$ transforms are enough to register n volumes whereas there are n unknown transforms τ^i in our problem. This results in an underdetermined problem: convergence is not guaranteed and the transforms can be subject to global drift. Following Wu et al. (2012), we remove the degree of freedom by adding an extra constraint to the control point displacements:

$$\mathcal{T}^* = \min_{\mathcal{T}} \mathbf{E}(\mathcal{T}) \quad s.t. \sum_{i=1}^n \mathbf{x}^i = \mathbf{0} \quad (9)$$

By traversing the set of matches \mathcal{M} , The energy function (8) can be rewritten as:

$$\mathbf{E}(\mathcal{T}) = \sum_{(p_a^i, p_b^j) \in \mathcal{M}} \left(s(p_a^i, p_b^j) d(p_a^i, p_b^j) \right)^2 \quad (10)$$

where $s(p_a^i, p_b^j)$ is defined by :

$$s(p_a^i, p_b^j) = \sqrt{w(p_a^i, p_b^j) \left(\frac{1}{|\mathcal{N}(p_a^i)|} + \frac{1}{|\mathcal{N}(p_b^j)|} \right)} \quad (11)$$

Equation (10) can be expressed in matrix notation :

$$\begin{aligned} \mathbf{A} &= \mathbf{S} \mathbf{M} (\mathbf{P} + \mathbf{B} \mathbf{X}) \\ \mathbf{E}(X) &= tr(\mathbf{A}^t \mathbf{A}) \end{aligned} \quad (12)$$

where \mathbf{S} is the $|\mathcal{M}| \times |\mathcal{M}|$ diagonal matrix of all $s(p_a^i, p_b^j)$.

\mathbf{P} is the $|\mathcal{P}| \times 3$ matrix of keypoint locations.

\mathbf{M} is the sparse $|\mathcal{M}| \times |\mathcal{P}|$ match matrix. Each line of \mathbf{M} represents a match between two keypoints, and contains exactly two nonzero values : one 1 and one -1 in the columns corresponding to the two matching keypoints. \mathbf{M} can be seen as the incidence matrix of the match directed graph where vertices are

Algorithm 1 One level of gradient descent for deformable registration. α is a small positive coefficient, $nIterations$ is the number of iterations.

```

 $\widetilde{\mathbf{X}}_0 \leftarrow 0$ 
 $k \leftarrow 0$ 
while  $k < nIterations$  do
  if  $k \bmod 10 = 0$  then
    Compute the parameters  $\theta^i$ 
  end if
  Compute  $\mathbf{S}_k$  from  $\theta^i$ 
   $\mathbf{X}_{k+1} \leftarrow \widetilde{\mathbf{X}}_k - \alpha 2\mathbf{B}^t \mathbf{M}^t \mathbf{S}_k^2 \mathbf{M} (\mathbf{P} + \mathbf{B}\widetilde{\mathbf{X}}_k)$ 
  Split  $\mathbf{X}_{k+1}$  into  $n$  coefficient matrices  $\mathbf{x}_{k+1}^i$ ,
   $\widetilde{\mathbf{x}}_{k+1}^i \leftarrow \mathbf{x}_{k+1}^i - \frac{1}{n} \sum_{j=1}^n \mathbf{x}_{k+1}^j$ 
  Construct  $\widetilde{\mathbf{X}}_{k+1}$  by stacking the matrices  $\widetilde{\mathbf{x}}_{k+1}^i$ 
   $k \leftarrow k + 1$ 
end while

```

keypoints and edges are matches with arbitrary orientation. An example of \mathbf{M} is shown by Figure 4.

\mathbf{B} is the $|\mathcal{P}| \times nc$ spline matrix. \mathbf{B} is a sparse non-square block diagonal matrix, where each non-zero block \mathbf{B}^i is the $|\mathcal{P}^i| \times c$ spline matrix for image i (the stacking of all $\mathbf{b}(p_a^i)$).

\mathbf{X} is the $nc \times 3$ stacking of all matrices \mathbf{x}^i .

To minimize equation (10), inspired by the Iteratively Reweighted Least Squares method, we use a gradient descent algorithm and we update the parameters θ^i every ten iterations. The gradient of equation (12) is :

$$\nabla \mathbf{E} = 2\mathbf{B}^t \mathbf{M}^t \mathbf{S}^2 \mathbf{M} (\mathbf{P} + \mathbf{B}\mathbf{X}) \quad (13)$$

After each iteration, we shift the mean of the spline coefficients in order to satisfy the constraint in (9). Note that an interesting byproduct of this step is the iterative removal of spatial bias in the common space (Guimond et al., 2000). Our gradient descent approach is summarized by algorithm 1.

We use a hierarchical approach, beginning with a coarse spline grid and incrementally doubling the grid density. This improves the speed of our approach and keeps the solution close to the global minimum, which is required as we use EM-weighting. Note that we achieve a significant simplification of the gradient computation (13), by taking into account the fact that the matrices \mathbf{M} and \mathbf{B} are sparse.

4.4. Guaranteed diffeomorphism

To guarantee that FROG outputs diffeomorphic half-transforms, we use a method similar to Rueckert et al. (2006): A sufficient condition to guarantee that a B-Spline based transformation is diffeomorphic is to check that for each direction (x, y and z), the displacement d of each control point is below $0.4 * g$

where g is the initial distance between control points (the grid spacing). Then, for a given resolution level, instead of using a single transformation which may not be diffeomorphic due to large control point displacements, we use a composition of several transformations, each respecting the diffeomorphism constraint. As the composition of diffeomorphisms is a diffeomorphism, the resulting composition is also a diffeomorphism.

For a given resolution level, optimization begins with a single grid for each half-transform. After each iteration, we check that the diffeomorphism constraint is fulfilled for all half-transforms. When the constraint is not satisfied, the iteration is canceled, current half-transforms are frozen, and iterations are resumed with a new set of half-transforms which will be composed with the previous ones. Subsequent iterations may also break the diffeomorphism condition, hence several transformations can be created for a single resolution level. As a result, for a given level, each half-transform is the concatenation of several transformations with limited displacements, which guarantees diffeomorphism. In our experiments on VISCERAL group A, with 3 levels of resolution, the number of concatenated transformations for each level was respectively 5, 8 and 14. Hence, instead of three grids (one for each level), our approach used a total of 27 grids per image. In our experiments, using a single grid for each resolution level yields half-transforms with a few percentage of voxels with negative Jacobian determinant, while using grid compositions yields no negative Jacobian determinant.

5. Results

5.1. Data sets and evaluation

We evaluate our approach with the VISCERAL database (Langs et al., 2013) which is composed of three groups of images (dimensions around $512 \times 512 \times 400$, spacing around $0.7 \times 0.7 \times 1.5\text{mm}$, 32 bits floats):

- A: 20 volumes (thorax and abdomen) of contrast enhanced (via contrast agent injection) CT scans. For these volumes, experts have located anatomical landmarks (up to 45 per volume, between 41 and 42 on average).
- B: 20 volumes (whole body) of CT scans. For these volumes, experts have located anatomical landmarks (up to 53 per volume, between 52 and 53 on average).
- C: 63 volumes (thorax and abdomen) of contrast enhanced CT scans, not annotated. The VISCERAL consortium refers to these volumes as the silver corpus.

With these three different groups (A, B and C), we are able to carry out several groupwise registrations with different scenarios: only A, only B, A and C or the three groups together, while observing the effects of these cases on the landmarks registration accuracy. To measure the quality of a given groupwise

registration, we first project all reference landmarks in the common space using the transforms τ^i . Then, for each landmark category (Clavicle L., Clavicle R. etc...), we compute a mean position \bar{p} . The quality criterion for this category is then defined as the average distance to \bar{p} . We also compute the global maximum individual distance to \bar{p} as a robustness criterion. Note that reference landmarks are only present for images in groups A and B. As a consequence, when registering groups A and C together, the criterion is restricted to images of group A. Note that landmarks are never used by any of the tested algorithm during registration. Experiments were carried out with a 24-core computer (48 logical cores with hyperthreading) with 128 GB of RAM.

We have compared our FROG approach with star-groupwise voxel-based methods NiftyReg (Modat et al., 2008, 2014) and ANTs (Avants et al., 2008). As the test data contains high resolution images, other dense methods could not be applied. As an example, Elastix (Klein et al., 2010) and GLIRT (Wu et al., 2012) require to store all images in-core, and GRAM (Hamm et al., 2010) exhibits a computational complexity which is too high for our data sets (see section 2.1).

5.2. Implementation and settings

We have implemented FROG in C++, and its code source is publicly available at <https://github.com/valette/frog>.

In the complexity analysis, the dominant terms are $|\mathcal{I}|v \log v$ for keypoint extraction, $|\mathcal{P}|^2$ for match computation and $|\mathcal{I}||\mathcal{P}|c \log c$ for optimization, where $|\mathcal{I}|$ is the number of volumes, v is the number of voxels for each volume, $|\mathcal{P}|$ is the total number of extracted keypoints, and c is the number of control points at the highest resolution.

Our pyramid of splines contains three levels, with a grid step of 200mm, 100mm, 50mm respectively. For each volume, we keep up to 20000 keypoints (see Agier et al. (2016)). For each resolution level, we compute 200 gradient descent iterations, with α set to 0.02.

5.3. Comparison with NiftyReg and ANTs

With NiftyReg, we have used the *groupwise_niftyreg_run.sh* script. This script computes a star-groupwise registration (see Figure 1(a)) and generates a template model from all input volumes. It alternates between mean reference computation and pairwise registration of all images with this reference. By default, 15 iterations of groupwise registration are applied: the first 5 iterations use rigid registration, the 10 last iterations use nonrigid registration.

With ANTs, we have used the *antsMultivariateTemplateConstruction2.sh* script. This script computes a star-groupwise registration (see Figure 1(a)) and generates a template model from all input volumes. It alternates between mean reference computation and pairwise registration of all images with this reference. By default, 5 iterations of groupwise registration are applied: the first is a rigid registration, the 4 others being nonrigid. The best results with ANTs have been obtained using mutual information as registration criterion.




Figure 6: Convergence of the minimized criterion and mean distance when registering group A. Distances are in mm. The blue curve is the square root of the criterion given by Equation (8) which reflects the average distance between matched keypoints. The black curve represents the global mean landmark distances. For the first 50 iterations, we use an anisotropic scale+translation model. For the next iterations, we use B-splines over 3 different resolution levels, 200 iterations for each level.

5.3.1. Group A

Detailed results are given by Tables 1 and 2.

ANTs processes the 20 volumes in 62.5 hours, with a mean registration distance of 15.5 mm. To further increase the speed of ANTs, we switched the registration subsampling factor sf from 1 to 2 and 4 (i.e. setting the shrinking factor parameter to 12x6x4x2 and 20x12x6x4, respectively). Setting $sf = 2$ increases both speed and registration quality: the 20 volumes are registered in 13.8 hours with a mean distance of 11.3 mm. In that case, there is no wrongly registered volume. Setting $sf = 4$ further increases computing speed but decreases output quality: the 20 volumes are registered in 9.4 hours with a mean distance of 14.0 mm. These results are presented in Table 2, where ANTS-S2 and ANTS-S4 refer to setting sf to 2 and 4, respectively.

NiftyReg registered group A in 67 hours on our test machine, yielding an average landmark distance of $9.6 \pm 7.5mm$.

FROG was able to successfully register all volumes of group A, while being much faster than NiftyReg and ANTs. Registration took 10mn on our test machine, and less than 30mn on a 4-core laptop. The two-tailed Welch’s t-test shows that our approach (Table 1 third column) achieves an average landmark distance ($8.7 \pm 7.2mm$) significantly lower to that of NiftyReg ($9.6 \pm 9.6mm$, $p - value = 0.006$) and ANTs ($11.3 \pm 7.9mm$, $p - value = 3.3 \cdot 10^{-18}$). We have also added experiments to table 2, varying the number of kept keypoints for each image to respectively 40k, 30k, 10k, 5k and 2.5k. Note that as expected, accuracy increases with the number of extracted keypoints, but saturates above 30k points per volume. Hence we have fixed the number of kept points to 20k. Figure 6 shows the convergence curve for \sqrt{E} (see equation 8) in blue and the mean landmarks distance during optimization in black. All distances are in mm.

Table 1 Registering Group A : comparison of landmark average distances (in mm) between our approach (FROG), NiftyReg and ANTs.

Landmark	ANTs	NiftyReg	FROG
Clavicle L	9.7 ± 4.6	6.8 ± 5.1	8.6 ± 2.4
Clavicle R	9.3 ± 4.4	7.1 ± 3.5	7.6 ± 3.4
Tubercul. L	11.2 ± 5.4	11.6 ± 6.1	9.3 ± 2.4
Tubercul. R	10.9 ± 5.5	11.7 ± 5.7	9.6 ± 5.2
C6	3.5 ± 2.6	4.1 ± 1.7	5.3 ± 2.3
C7	5.0 ± 3.3	7.6 ± 3.8	6.7 ± 4.3
Th1	4.5 ± 3.0	7.6 ± 4.2	6.7 ± 5.6
Th2	4.9 ± 3.4	6.6 ± 3.7	6.0 ± 4.9
Th3	6.1 ± 4.2	6.9 ± 3.8	5.7 ± 4.4
Th4	7.6 ± 4.3	7.9 ± 3.4	5.9 ± 4.6
Th5	8.6 ± 3.6	8.4 ± 3.0	6.8 ± 4.6
Th6	9.6 ± 3.4	9.9 ± 3.5	7.8 ± 5.1
Th7	10.2 ± 3.4	10.2 ± 3.8	8.4 ± 6.1
Th8	10.3 ± 3.9	11.0 ± 4.7	8.9 ± 8.2
Th9	10.8 ± 4.9	10.8 ± 3.8	10.8 ± 9.8
Th10	11.3 ± 5.4	11.3 ± 3.6	11.7 ± 11.1
Th11	12.2 ± 5.9	11.8 ± 5.9	13.0 ± 10.3
Th12	13.4 ± 6.2	11.8 ± 7.3	13.5 ± 9.6
L1	13.3 ± 6.6	11.9 ± 7.6	12.7 ± 9.5
L2	12.8 ± 6.6	11.6 ± 8.3	10.6 ± 7.8
L3	13.1 ± 7.1	11.0 ± 8.3	11.6 ± 6.5
L4	12.8 ± 7.2	10.0 ± 7.2	10.1 ± 6.0
L5	11.4 ± 6.7	8.4 ± 6.7	7.8 ± 6.4
Sternoclav. L	7.9 ± 5.5	8.5 ± 5.7	5.4 ± 4.7
Sternoclav. R	6.6 ± 3.9	5.0 ± 4.4	4.9 ± 3.3
Aortic arch	8.4 ± 3.2	9.3 ± 3.7	9.1 ± 3.9
Trachea bif.	4.5 ± 3.8	4.5 ± 4.0	4.6 ± 2.0
Bronchus L	7.5 ± 5.5	8.2 ± 8.6	6.8 ± 3.6
Bronchus R	5.9 ± 4.0	6.4 ± 7.8	4.5 ± 2.0
Coronaria	9.1 ± 6.4	10.1 ± 10.3	7.5 ± 2.6
Aortic valve	11.8 ± 7.2	13.7 ± 10.1	9.4 ± 3.8
Xyphoideus	16.6 ± 7.8	15.5 ± 9.6	13.5 ± 9.8
Renal pelvis L	17.9 ± 9.2	10.9 ± 10.6	7.3 ± 3.3
Renal pelvis R	20.6 ± 15.1	16.0 ± 13.4	11.9 ± 13.3
Crista iliaca L	9.9 ± 6.0	9.4 ± 6.1	9.7 ± 6.0
Crista iliaca R	10.4 ± 8.8	10.4 ± 8.7	9.5 ± 7.3
Aorta bif.	15.6 ± 10.5	11.2 ± 7.3	9.4 ± 5.4
VCI bif.	12.6 ± 6.5	9.1 ± 5.2	8.2 ± 4.3
Troch. maj. L	19.3 ± 10.8	15.8 ± 9.1	15.9 ± 8.8
Troch. maj. R	18.5 ± 9.1	16.5 ± 9.7	14.9 ± 8.4
Ischiadicum L	11.8 ± 6.4	4.3 ± 3.5	4.6 ± 3.2
Ischiadicum R	10.9 ± 6.3	4.8 ± 2.9	5.7 ± 3.0
Symphysis	15.8 ± 9.9	10.7 ± 8.5	10.2 ± 8.0
Troch. min. L	15.2 ± 7.7	3.7 ± 2.0	5.3 ± 1.9
Troch. min. R	16.6 ± 5.8	4.1 ± 2.7	6.1 ± 3.0
Mean	11.3 ± 7.9	9.6 ± 7.5	8.7 ± 7.2
Maximum	70.6	57.5	65.4
Time (h)	13.8	67.0	0.2




Figure 7: Comparison between our FROG approach and NiftyReg on VISCERAL group A: landmarks transformed in the common space.

Figure 7 compares reference landmarks transformed in the common space, for NiftyReg and our approach. Figure 8 summarizes the statistics with boxplots of landmark errors obtained with ANTs (group A), NiftyReg (group A), FROG (group A) and FROG (groups A+C). On each plot, the central mark is the median, the edges of the box are the 25th (Q1) and 75th percentiles (Q3), the whiskers ($W = 1.5$) extend to the most extreme data points not considered to be outliers, and the outliers are plotted individually if they are larger than $Q3 + W(Q3 - Q1)$. $W = 1.5$ corresponds approximately to ± 2.7 standard deviation and 99.3 coverage if the data is normally distributed. Notches draw comparison intervals. Two medians are significantly different at the 5% level if their intervals do not overlap. The interval endpoints are the extremes of the notches. If the notches in the box plot do not overlap, we can conclude with 95% confidence that the true medians do differ.

5.3.2. Group B

Figure 9 compares groupwise registration of group B. We have extracted bone structures from each volume and superposed them in the common space, using (a) ANTs, (b) NiftyReg and (c) FROG. NiftyReg yields a better result on




Figure 8: Landmark distances on the set of reference landmarks transformed into the common space. From left to Right : ANTs (group A), NiftyReg (group A), FROG (group A), FROG (groups A+C). Distances are in mm.

the torso and hip regions, but still lacks robustness for the legs. Our approach is much more robust to variability, as most of its error is concentrated around the patient arms, which are very challenging regions due to different shoulder poses. The average landmark distance (reported in the second block of Table 2) obtained with our approach is $9.2 \pm 8.3mm$, significantly lower than NiftyReg ($12.8 \pm 11.6mm$) and ANTs ($17.9 \pm 27.3mm$). Note that landmark distances are larger than the ones obtained on group A, for all algorithms (Ours, NiftyReg, ANTs). One reason is that the images of group B were obtained without any contrast agent. As a consequence, soft tissue is much less visible on group B, which could explain lower accuracy.

5.3.3. Groups A and C

When registering groups A and C together, ANTs yields a mean distance of 44.7 mm. One explanation is that computing the average image from 83 full body scans with a high variability can lead to a very blurred image when registration is not perfect (top right image in Figure 10). As a consequence, registration against a blurry reference image degenerates.

Registering groups A and C together took 324 hours with NiftyReg and has resulted in an average landmark distance of $9.3 \pm 9.3mm$.

Our algorithm registers groups A and C within 1.3 hours and yields an average distance of $8.5 \pm 6.9mm$, statistically better than NiftyReg ($9.3 \pm 9.3mm$, $p - value = 0.0045$) and significantly better than ANTs ($44.7 \pm 31.6mm$). We have estimated a rate of outliers of 70% when registering groups A and C together. Note that our previous work based on rigid groupwise registration (Agier et al., 2016) yields an average error of 31.7mm.

Table 2 Comparison between our approach, ANTs and NiftyReg on the VIS-CERAL data set. Our approach extracts 20k points per volume, except for lines 2 to 6 where the number of points vary between 40k and 2.5k

Groups	N	Algorithm	Max d (mm)	\bar{d} (mm)	Time (h)
A	20	FROG	65.4	8.7 ± 7.2	0.2
A	20	FROG (40k)	67.5	8.7 ± 7.1	0.4
A	20	FROG (30k)	65.9	8.7 ± 7.1	0.2
A	20	FROG (10k)	71.6	9.2 ± 7.6	0.1
A	20	FROG (5k)	63.9	9.9 ± 8.1	0.1
A	20	FROG (2.5k)	68.9	10.9 ± 8.7	0.1
A	20	FROG (SIFT)	70.9	11.6 ± 9.7	4.2
A	20	ANTs-S4	152.7	14.0 ± 17.0	9.4
A	20	ANTs-S2	70.6	11.3 ± 7.9	13.8
A	20	ANTs	150.2	15.5 ± 16.7	62.5
A	20	NiftyReg	57.5	9.6 ± 7.5	67.0
B	20	FROG	64.9	9.2 ± 8.3	0.5
B	20	ANTs-S2	231.0	17.9 ± 27.3	36.0
B	20	NiftyReg	105.0	12.8 ± 11.6	135.0
AC	83	FROG	60.4	8.5 ± 6.9	1.3
AC	83	ANTs-S2	198.0	44.7 ± 31.6	82.8
AC	83	NiftyReg	64.0	9.3 ± 9.3	324.0
ABC	103	FROG	61.7	8.7 ± 7.0	1.9

5.3.4. Overall

Table 3 shows timings for our algorithm as well as the average number of matches per volume for several scenarios. Note that even if the number of keypoints per volume, is constant (20k), the average number of matches increases with the number of volumes. As an example, when registering group A, each keypoint is matched with $190k/20k = 9.5$ points on average. But when registering groups A and C together, each keypoint is matched with $820k/20k = 41$ points on average. In terms of memory footprint, our approach and NiftyReg never needed more than 10 GB, while we observed a peak memory consumption greater than 65 GB with ANTs at full resolution, and 34 GB when setting $sf = 2$.

Finally, figure 10 shows the average of the registered images. From top to bottom: ANTs, NiftyReg and our algorithm, when registering Group A (left) and groups A and C (right). NiftyReg and ANTs provide a sharper image than our approach when registering group A. When registering groups A and C (83 images), the resulting average image is blurred for ANTs, whereas our approach remains robust to the number of images. NiftyReg still yields an average image sharper than ours. Hence, our approach is not very effective for template construction (average image). But, in contrast with NiftyReg and ANTs, our hub-less approach does not need any average image during registration. Our average image is only a byproduct used for comparison, computed after the registration. Moreover, image-based criteria have been shown unreliable to evaluate registra-

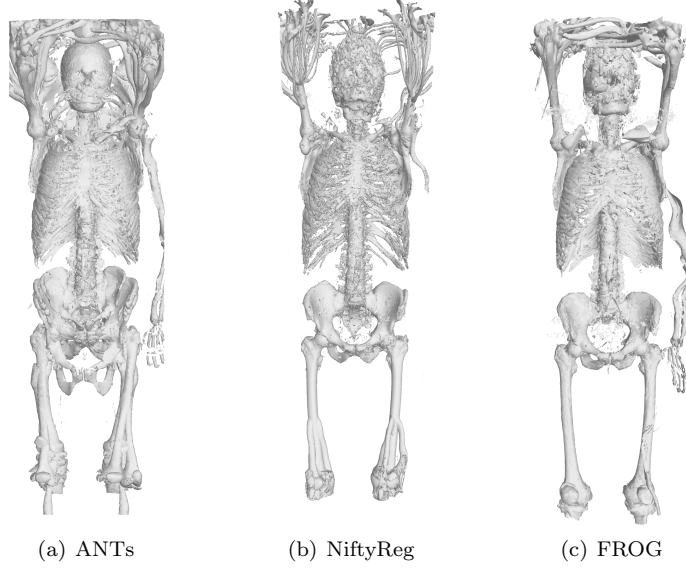


Figure 9: Bone structures extracted from a registered group of 20 whole-body CT volumes of the VISCERAL group B. Groupwise registration has been carried out with (a) ANTs, (b) NiftyReg and (c) our algorithm. Our algorithm exhibits a better overall robustness to variability.

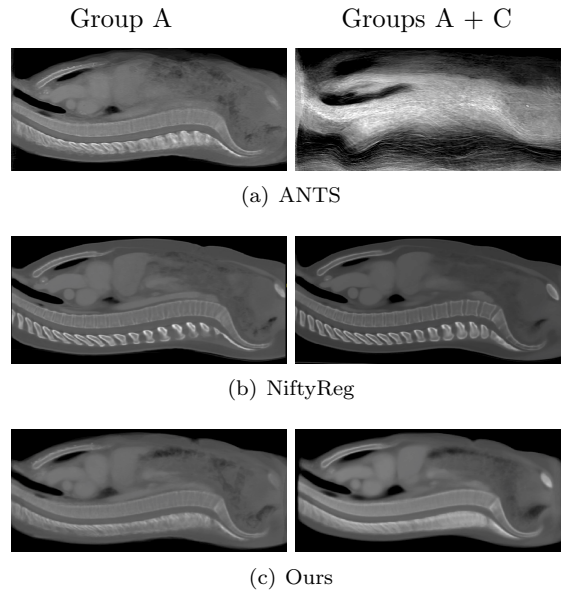


Figure 10: Average images obtained with (a) ANTs, (b) NiftyReg, (c) Ours. Left : when registering group A. Right : when registering groups A+C.

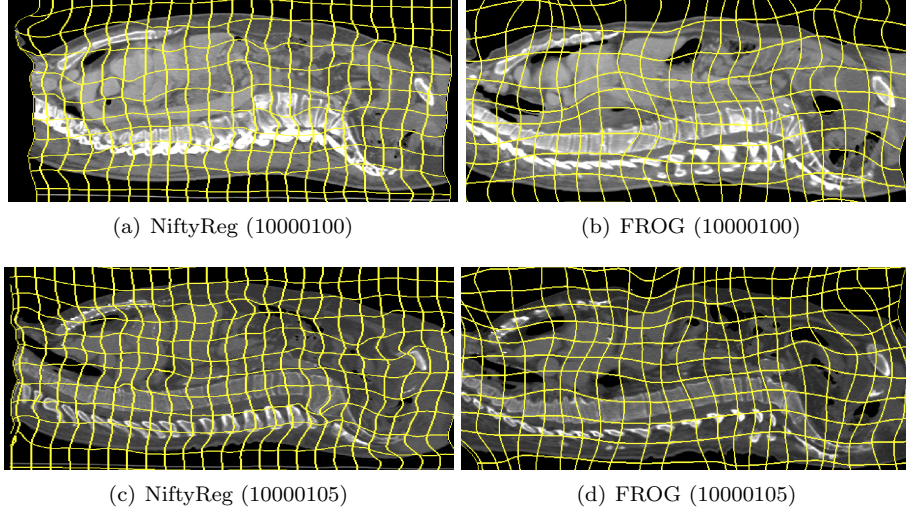


Figure 11: Comparison of transformed images for two individuals (10000100 and 10000105 in VISCERAL Group A): Despite providing a sharper mean image (see figure 10) NiftyReg sometimes yields transformations that are less realistic than FROG. This can be easily observed on some oversized lumbar vertebrae in 10000100. FROG also sometimes yields unrealistic deformation, such as the chest boundary in 10000105.

tion accuracy (Rohlfing, 2012). We observed that even though NiftyReg yields a sharp mean image, individual transforms can be unrealistic. As an example, figure 11 compares the images 1000100 and 1000105 from VISCERAL group A transformed by NiftyReg and FROG. The overlay grid reflects the deformation for all transforms. Some lumbar vertebrae are clearly oversized with NiftyReg on image 1000100. On the other hand, FROG sometimes also deforms images in an unrealistic way, as shown on the chest boundary in image 1000105. Finally, although our average image is not the sharpest, landmark distances (table 1) are on average lower with our approach.

We have also registered all three groups A, B, and C in less than 2 hours, with a reasonable error, which illustrates the ability of our approach to register more than 100 images, mixing images obtained with or without contrast agent.

5.4. Using 3D-SIFT

We have carried out an experiment with points defined with 3D-SIFT (Rister et al., 2017), visible on line 7 in table 2. Keypoint extraction is much slower than our 3D-SURF approach, which results in a registration time of 4.2 hours, while yielding larger landmark distances (11.6mm for 3D-SIFT vs 9.0mm for 3D-SURF). This shows that our approach is able to exploit various keypoint types, and also illustrates the relevance of choosing 3D-SURF over 3D-SIFT, as reported initially by Bay et al. (2006).

Table 3 Number of matches per volume and processing times for our algorithm. n is the number of registered volumes, $|\mathcal{M}|/n$ the number of matches per volume. Columns I, II and III indicate the time spent (in minutes) for keypoint extraction, matching and optimization, respectively.

	Data	n	$ \mathcal{M} /n$	I (mn)	II (mn)	III (mn)	Total (mn)
	A	20	190k	6.3	1.5	2.2	10.0
	B	20	190k	24.4	1.5	2.6	28.5
	AC	83	820k	30.2	28.5	20.6	79.3
	ABC	103	1019k	52.8	31.3	32.8	116.9

5.5. Limitation

The limitation of our proposal is related to the requirement of a sufficient number of inlier keypoints in each cell of the grid at the highest resolution. In other words the size of the finest grid used is locally limited by the keypoint density. This results in a limit in registration accuracy. We observed this limitation in processing the MRI LONI Probabilistic Brain Atlas (LPBA40) (Shattuck et al., 2008), containing 40 human brain images of resolution $256 \times 124 \times 256$. Processing time for the whole dataset is less than 10 minutes. The resulting average image is displayed in Figure 12. The accuracy of our generic approach is not sufficient to correctly handle small structures inside the brain. Hence, our approach is not as accurate as algorithms dedicated to brain processing such as GLIRT (Wu et al., 2012), which reports an average overlap ratio of 67% (about 20% higher than ours). Yet our algorithm converges in a short time, which illustrates the versatility of our approach. In practice, our approach is able to extract about 10k keypoints for each brain volume of LPBA40, while more than 100k points have been extracted from each volume of the VISCERAL database. Thus, for the VISCERAL dataset, we retained only the most significant points, discarding weak points. This was not possible for the LPBA40 database. This comes from the fact that brain MRI images are less sharp and contain fewer voxels than full body CT scans. Hence, a possible improvement could be to develop a keypoint extractor more adapted to brain MRI, which could yield much more keypoints and therefore increase registration accuracy.

6. Conclusion

Our algorithm is able to register CT and MRI volumes in reasonable time and low memory consumption. Experimental results illustrate its robustness to variability, and comparisons with NiftyReg and ANTs show that our algorithm yields lower mean landmark distances within a much shorter computation time. Our approach is currently suited for large databases of relatively high resolution images. As an example, brain databases generally contain images with $256 \times 256 \times 256$ voxels, but our approach is not as accurate as state-of-the-art voxel-based approaches for these resolutions. A possible improvement could be a keypoint extractor dedicated to MRI images with fewer voxels and fewer sharp




Figure 12: Average brain when registering the LPBA40 MRI brain dataset with our approach.

features. But many applications could already benefit from our approach. As an example, the low complexity and memory requirement of our algorithm allows to use a full groupwise nonrigid registration algorithm as a preliminary task for applications such as atlas based segmentation, machine learning classification and longitudinal studies. A preliminary application using FROG is proposed in Franchi et al. (2019), where we predict patient gender from 3D CT images using groupwise registration.

As depicted in figure 1, we use a full graph approach. As a consequence, the number of matches grows faster than the number of images. Hence, the number of matches per images increases with the number of images (see Table 3). While adding more images has the potential to produce more matching keypoints to improve registration accuracy, our results did not demonstrate statistically significant improvement.

Our method is oblivious to point cloud sources and is not restricted to SURF keypoints and splines. Hence, investigating keypoint descriptors tailored to medical images and richer transform models are two very interesting research perspectives for our work. Computing features with machine learning seems a promising option in this context.

As our algorithm provides a robust approach able to handle keypoints between different patients without needing dense voxel information during optimization, it paves the way towards fast algorithms able to deal with large data sets. We showed that although our approach exhibits quadratic complexity and memory footprint, it is able to register more than one hundred images in a short time (less than 2 hours). With our current hardware (128GB RAM), we project to be able to register as many as 1000 volumes with FROG, in about 125 hours. Note that this would still be less than the time taken by NiftyReg to register 83 volumes. Scaling the algorithm to even larger data sets (more than 10000 patients) is still a challenging problem.

Acknowledgment

Many thanks to the VISCERAL Consortium (del Toro et al., 2016) (www.visceral.eu) for allowing us to compute evaluations using the VISCERAL data set.

References

- Agier, R., Valette, S., Fanton, L., Croisille, P., Prost, R., 2016. Hubless 3D medical image bundle registration, in: VISAPP 2016 11th Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, pp. 265–272.
- Allaire, S., Kim, J.J., Breen, S.L., Jaffray, D.A., Pekar, V., 2008. Full orientation invariance and improved feature selectivity of 3D SIFT with application to medical image analysis, in: Computer Vision and Pattern Recognition Workshops, 2008. CVPRW’08. IEEE Computer Society Conference on, IEEE. pp. 1–8.
- Anguelov, D., Dulong, C., Filip, D., Frueh, C., Lafon, S., Lyon, R., Ogale, A., Vincent, L., Weaver, J., 2010. Google street view: Capturing the world at street level. *Computer* 43.
- Ashburner, J., Friston, K.J., Frackowiak, R., 2003. Rigid body registration. San Diego, CA: Academic Press.
- Avants, B., Epstein, C., Grossman, M., Gee, J., 2008. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis* 12, 26 – 41.
- Bartoli, A., Pizarro, D., Loog, M., 2013. Stratified generalized procrustes analysis. *International Journal of Computer Vision* 101, 227–253.
- Bay, H., Tuytelaars, T., Van Gool, L., 2006. SURF: Speeded up robust features, in: Computer Vision–ECCV 2006. Springer, pp. 404–417.
- Brown, L.G., 1992. A survey of image registration techniques. *ACM Comput. Surv.* 24, 325–376. doi:10.1145/146370.146374.
- Cheung, W., Hamarneh, G., 2009. n-SIFT: n-dimensional scale invariant feature transform. *IEEE Trans. Image Processing* 18, 2012–2021. doi:10.1109/TIP.2009.2024578.
- Chui, H., Rangarajan, A., 2003. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding* 89, 114 – 141. URL: <http://www.sciencedirect.com/science/article/pii/S1077314203000092>, doi:[https://doi.org/10.1016/S1077-3142\(03\)00009-2](https://doi.org/10.1016/S1077-3142(03)00009-2). nonrigid Image Registration.

- Datta, R., Li, J., Wang, J.Z., 2005. Content-based image retrieval: approaches and trends of the new age, in: *Proceedings of the 7th ACM SIGMM International Workshop on Multimedia Information Retrieval*, ACM. pp. 253–262.
- Dauphin, Y.N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., Bengio, Y., 2014. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization, in: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (Eds.), *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., pp. 2933–2941.
- Dempster, A.P., Laird, N.M., Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* , 1–38.
- Donoho, D.L., Huber, P.J., 1983. The notion of breakdown point. *A festschrift for Erich L. Lehmann* 157184.
- Fonov, V., Evans, A.C., Botteron, K., Almli, C.R., McKinstry, R.C., Collins, D.L., 2011. Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage* 54, 313 – 327.
- Fornet, M., Rohr, K., Stiehl, H.S., 2001. Radial basis functions with compact support for elastic registration of medical images. *Image and Vision Computing* 19, 87–96.
- Fox, J., 2002. Robust regression. *An R and S-Plus Companion to Applied Regression* .
- Frahm, J.M., Fite-Georgel, P., Gallup, D., Johnson, T., Raguram, R., Wu, C., Jen, Y.H., Dunn, E., Clipp, B., Lazebnik, S., et al., 2010. Building rome on a cloudless day, in: *Computer Vision–ECCV 2010*. Springer, pp. 368–381.
- Franchi, A., Valette, S., Agier, R., Prost, R., Kéchichan, R., Fanton, L., 2019. The prospects for application of computational anatomy in forensic anthropology for sex determination. *Forensic Science International* 297, 156 – 160.
- Guimond, A., Meunier, J., Thirion, J.P., 2000. Average brain models: A convergence study. *Computer Vision and Image Understanding* 77, 192 – 210.
- Hamm, J., Ye, D.H., Verma, R., Davatzikos, C., 2010. GRAM: A framework for geodesic registration on anatomical manifolds. *Medical Image Analysis* 14, 633–42.
- Harris, C., Stephens, M., 1988. A combined corner and edge detector., in: *Alvey Vision Conference*, Manchester, UK. p. 50.
- Hartley, R., Zisserman, A., 2003. Multiple view geometry in computer vision. Cambridge university press.

- Huang, B., Wang, W., Bates, M., Zhuang, X., 2008. Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy. *Science* 319, 810–813.
- Huber, P.J., 2011. Robust statistics. Springer.
- Iglesias, J.E., Sabuncu, M.R., 2015. Multi-atlas segmentation of biomedical images: A survey. *Medical Image Analysis* 24, 205 – 219.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17, 825–841.
- Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23, S151–S160.
- Karlsson, N., Di Bernardo, E., Ostrowski, J., Goncalves, L., Pirjanian, P., Munich, M.E., 2005. The vSLAM algorithm for robust localization and mapping, in: *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, IEEE. pp. 24–29.
- Keall, P.J., Joshi, S., Vedam, S.S., Siebers, J.V., Kini, V.R., Mohan, R., 2005. Four-dimensional radiotherapy planning for dmlc-based respiratory motion tracking. *Medical Physics* 32, 942–951.
- Kéchichian, R., Valette, S., Sdika, M., Desvignes, M., 2014. Automatic 3d multiorgan segmentation via clustering and graph cut using spatial relations and hierarchically-registered atlases, in: Menze, B., Langs, G., Montillo, A., Kelm, M., Müller, H., Zhang, S., Cai, W., Metaxas, D. (Eds.), *Medical Computer Vision: Algorithms for Big Data*, Springer International Publishing. Springer International Publishing. pp. 201–209.
- Klein, S., Staring, M., Murphy, K., Viergever, M., Pluim, J., 2010. Elastix: a toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imag.* 29, 196 – 205.
- Läbe, T., Förstner, W., 2004. Geometric stability of low-cost digital consumer cameras, in: *Proceedings of the 20th ISPRS Congress*, Istanbul, Turkey, pp. 528–535.
- Langs, G., Hanbury, A., Menze, B., Müller, H., 2013. VISCERAL: Towards large data in medical imaging – challenges and directions, in: *Proceedings of the Third MICCAI International Conference on Medical Content-Based Retrieval for Clinical Decision Support*, Springer-Verlag, Berlin, Heidelberg. pp. 92–98. doi:10.1007/978-3-642-36678-9_9.
- Lewis, J., 1995. Fast normalized cross-correlation, in: *Vision Interface*, pp. 120–123.

- Li, B., Christensen, G.E., Hoffman, E.A., McLennan, G., Reinhardt, J.M., 2012. Establishing a normative atlas of the human lung. *Academic Radiology* 19, 1368–1381.
- López, A.M., Lumbreras, F., Serrat, J., Villanueva, J.J., 1999. Evaluation of methods for ridge and valley detection. *IEEE Trans. Pattern Anal. Machine Intell.* 21, 327–335.
- Lötjönen, J.M., Wolz, R., Koikkalainen, J.R., Thurfjell, L., Waldemar, G., Soininen, H., Rueckert, D., Initiative, A.D.N., et al., 2010. Fast and robust multi-atlas segmentation of brain magnetic resonance images. *NeuroImage* 49, 2352–2365.
- Lowe, D.G., 1999. Object recognition from local scale-invariant features, in: *Computer Vision, proceedings of the seventh IEEE International Conference on*, pp. 1150–1157.
- Lowe, D.G., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 91–110.
- Maintz, J., Viergever, M.A., 1998. A survey of medical image registration. *Medical Image Analysis* 2, 1–36.
- Michalski, R.S., Carbonell, J.G., Mitchell, T.M., 2013. *Machine learning: An artificial intelligence approach*. Springer Science & Business Media.
- Mikolajczyk, K., Schmid, C., 2002. An affine invariant interest point detector, in: *European Conference on Computer Vision*, Springer. pp. 128–142.
- Mikolajczyk, K., Schmid, C., 2005. A performance evaluation of local descriptors. *IEEE Trans. Pattern Anal. Machine Intell.* 27, 1615–1630.
- Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S., 2014. Global image registration using a symmetric block-matching approach. *Journal of Medical Imaging* 1, 1 – 1 – 6.
- Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., 2008. Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine* 98, 278–284.
- Moré, J.J., 1978. The Levenberg-Marquardt algorithm: implementation and theory, in: *Numerical Analysis*. Springer, pp. 105–116.
- Murail, P., Bruzek, J., Houet, F., Cunha, E., 2008. DSP: A tool for probabilistic sex diagnosis using worldwide variability in hip-bone measurements. *Bulletins et mémoires de la Société d’Anthropologie de Paris* 17.
- Myronenko, A., Song, X., 2010. Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 2262–2275. doi:10.1109/TPAMI.2010.46.

- Pluim, J.P., Maintz, J.A., Viergever, M.A., 2003. Mutual-information-based registration of medical images: a survey. *IEEE Trans. Med. Imag.* 22, 986–1004.
- Rangarajan, A., Chui, H., Mjolsness, E., Pappu, S., Davachi, L., Goldman-Rakic, P., Duncan, J., 1997. A robust point-matching algorithm for autoradiograph alignment. *Medical Image Analysis* 1, 379 – 398. URL: <http://www.sciencedirect.com/science/article/pii/S1361841597850086>, doi:[https://doi.org/10.1016/S1361-8415\(97\)85008-6](https://doi.org/10.1016/S1361-8415(97)85008-6).
- Rasoulouian, A., Rohling, R., Abolmaesumi, P., 2012. Group-wise registration of point sets for statistical shape models. *IEEE Trans. Med. Imag.* 31, 2025–2034.
- Rister, B., Horowitz, M.A., Rubin, D.L., 2017. Volumetric image registration from invariant keypoints. *IEEE Trans. Image Processing* 26, 4900–4910. doi:10.1109/TIP.2017.2722689.
- Robinson, D., Milanfar, P., 2004. Fundamental performance limits in image registration. *IEEE Trans. Image Processing* 13, 1185–1199.
- Rohlfing, T., 2012. Image similarity and tissue overlaps as surrogates for image registration accuracy: Widely used but unreliable. *IEEE Trans. Med. Imag.* 31, 153–163.
- Rueckert, D., Aljabar, P., Heckemann, R.A., Hajnal, J.V., Hammers, A., 2006. Diffeomorphic registration using b-splines, in: Larsen, R., Nielsen, M., Sparring, J. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 702–709.
- Scahill, R.I., Frost, C., Jenkins, R., Whitwell, J.L., Rossor, M.N., Fox, N.C., 2003. A longitudinal study of brain volume changes in normal aging using serial registered magnetic resonance imaging. *Archives of Neurology* 60, 989–994.
- Schönemann, P.H., 1966. A generalized solution of the orthogonal procrustes problem. *Psychometrika* 31, 1–10.
- Shattuck, D.W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K.L., Poldrack, R.A., Bilder, R.M., Toga, A.W., 2008. Construction of a 3D probabilistic atlas of human cortical structures. *NeuroImage* 39, 1064 – 1080.
- Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable medical image registration: A survey. *IEEE Trans. Med. Imag.* 32, 1153–1190.
- Suh, J.W., Kwon, O.K., Scheinost, D., Sinusas, A.J., Cline, G.W., Papademetris, X., 2011. Whole body nonrigid ct-pet registration using weighted demons, in: *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 1223–1226. doi:10.1109/ISBI.2011.5872622.

- Szeliski, R., Coughlan, J., 1997. Spline-based image registration. *International Journal of Computer Vision* 22, 199–218.
- del Toro, O.J., Muller, H., Krenn, M., Gruenberg, K., Taha, A.A., Winterstein, M., Eggel, I., Foncubierta-Rodriguez, A., Goksel, O., Jakab, A., Kontokotsios, G., Langs, G., Menze, B., Fernandez, T.S., Schaer, R., Walleyo, A., Weber, M.A., Cid, Y.D., Gass, T., Heinrich, M., Jia, F., Kahl, F., Kechichian, R., Mai, D., Spanier, A., Vincent, G., Wang, C., Wyeth, D., Hanbury, A., 2016. Cloud-based evaluation of anatomical structure segmentation and landmark detection algorithms: VISCERAL anatomy benchmarks. *IEEE Trans. Med. Imag.* 35, 2459–2475. doi:10.1109/TMI.2016.2578680.
- Triggs, B., McLauchlan, P.F., Hartley, R.I., Fitzgibbon, A.W., 2000. Bundle adjustment – a modern synthesis, in: *Vision Algorithms: Theory and Practice*. Springer, pp. 298–372.
- Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2007. Non-parametric diffeomorphic image registration with the demons algorithm, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer. pp. 319–326.
- Wachinger, C., Toews, M., Langs, G., Wells, W., Golland, P., 2018. Keypoint transfer for fast whole-body segmentation. *IEEE Transactions on Medical Imaging*, to appear.
- Wang, F., Vemuri, B.C., Rangarajan, A., 2006. Groupwise point pattern registration using a novel CDF-based Jensen-Shannon Divergence, in: *Computer Vision and Pattern Recognition (CVPR), 2006 IEEE Conference on*, pp. 1283–1288.
- Wang, H., Kläser, A., Schmid, C., Liu, C.L., 2011. Action recognition by dense trajectories, in: *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, IEEE. pp. 3169–3176.
- Wu, G., Jia, H., Wang, Q., Shen, D., 2011. Sharpmean: groupwise registration guided by sharp mean image and tree-based registration. *NeuroImage* 56, 1968–1981.
- Wu, G., Wang, Q., Jia, H., Shen, D., 2012. Feature-based groupwise registration by hierarchical anatomical correspondence detection. *Human Brain Mapping* 33, 253–271.
- Xu, Z., Lee, C.P., Heinrich, M.P., Modat, M., Rueckert, D., Ourselin, S., Abramson, R.G., Landman, B.A., 2016. Evaluation of six registration methods for the human abdomen on clinically acquired ct. *IEEE Transactions on Biomedical Engineering* 63, 1563–1572. doi:10.1109/TBME.2016.2574816.
- Ying, S., Wu, G., Wang, Q., Shen, D., 2014. Hierarchical unbiased graph shrinkage (HUGS): a novel groupwise registration for large data set. *NeuroImage* 84, 626–638.

- Zhang, P., Cootes, T.F., 2012. Automatic construction of parts+geometry models for initializing groupwise registration. *IEEE Trans. Med. Imag.* 31, 341–358.
- Zheng, X., Zhou, M., Wang, X., 2008. Interest point based medical image retrieval, in: *Medical Imaging and Informatics*. Springer, pp. 118–124.
- Zhu, D., Li, K., Guo, L., Jiang, X., Zhang, T., Zhang, D., Chen, H., Deng, F., Faraco, C., Jin, C., Wee, C.Y., Yuan, Y., Lv, P., Yin, Y., Hu, X., Duan, L., Hu, X., Han, J., Wang, L., Shen, D., Miller, L.S., Li, L., Liu, T., 2013. DICCOL: Dense individualized and common connectivity-based cortical landmarks. *Cerebral Cortex* 23, 786–800.

