

Deformable Mosaicing for Whole-body MRI^{*}

Christian Wachinger¹, Ben Glocker^{1,3}, Jochen Zeltner², Nikos Paragios³,
Nikos Komodakis⁴, Michael Sass Hansen⁵, and Nassir Navab¹

¹ Computer Aided Medical Procedures (CAMP), TUM, Munich, Germany
wachinge@cs.tum.edu

² Siemens Medical Solutions, Erlangen, Germany

³ Mathématiques Appliquées aux Systèmes (MAS), Ecole Centrale Paris, France

⁴ Computer Science Department, University of Crete, Greece

⁵ Technical University of Denmark

Abstract. Whole-body magnetic resonance imaging is an emerging application gaining vast clinical interest during the last years. Although recent technological advances shortened the longish acquisition time, this is still the limiting factor avoiding its wide-spread clinical usage. The acquisition of images with large field-of-view helps to relieve this drawback, but leads to significantly distorted images. Therefore, we propose a *deformable mosaicing* approach, based on the simultaneous registration to linear weighted averages, to correct for distortions in the overlapping area. This method produces good results on in-vivo data and has the advantage that a seamless integration into the clinical workflow is possible.

1 Introduction

Whole-body (WB) magnetic resonance imaging (MRI) is becoming a popular clinical tool due to the recent technological advances in MRI, making faster acquisitions possible. Unlike computed tomography (CT), the acquisition of high-resolution MR images is not feasible during continuous table movement, making a multi-station scanning necessary to cover larger body regions. The compounding of the partially overlapping volumes is straightforward, since the MR scanner keeps track of their exact spatial locations.

The creation of WB images further increases the number of clinical applications for MRI, so far reserved for other modalities, see section 2. From a current perspective, the major disadvantage using MRI for WB imaging in comparison to CT is the longer scanning time. In this report, we use MR acquisitions with a large field-of-view (FOV), enabling to cover with the same number of scans larger parts of the body. This, however, leads to a degradation of the images by geometrical distortion artefacts towards the boundaries, further described in section 3. We propose a novel method, originating from the field of brain atlas construction, to correct for the geometrical distortion in the overlapping area, see section 4. Our experiments show the good results on in-vivo data, see section 5.

^{*} This work is partially supported by Siemens Medical Solutions, Germany.

2 Clinical Applications of WB-MRI

Recent advances in MRI such as multi-channel receiver, parallel imaging techniques, and automated table movement make high-resolution WB-MRI clinically feasible [1].

First clinical studies show its value in oncological applications, which focus on the search for metastases of cancer patients in the whole body. Lauenstein *et al.* [2] found out that WB-MRI compares well with standard methods such as CT, PET, and nuclear scintigraphy for the detection of cerebral, pulmonary, and hepatic lesions and is more sensitive for the detection of hepatic and osseous metastases. Additionally, the excellent contrast of MRI provides further information about soft tissue and organs to the physician, supporting his diagnosis [3].

Non-oncological applications are whole-body fat measurement to evaluate body composition and muscular infections, angiography for the diagnosis of atherosclerosis, and virtual autopsy in forensic science [3]. WB-MRI is also the method of choice for screening and prevention purposes, which is ethically questionable with CT due to radiation exposure of healthy persons. Moreover, contrast agents used for MRI to highlight specific organs or the vascular tree are relatively safe in comparison to iodine based ones used for CT [3].

These reasons indicate that MRI is challenging CT as standard WB imaging modality. For head and abdomen, the superior performance of MR was already noted, but imaging the lungs is still delicate [3]. The major drawback, however, remains the longer acquisition time with MRI. We address this issue by using scans with a larger FOV, allowing for covering an equivalent region with less stations. They have the same resolution as scans with a regular FOV, leading to the same acquisition time, but an increased voxel spacing, leading to a negligible loss of image quality. Typical is a normal FOV of 35 cm and an enlarged one of 50 cm. A disadvantage of the enlarged FOV is that spins are excited, which are farther away from the magnetic iso-center and therefore more sensitive to geometrical distortions, see section 3.




Fig. 1: WB-MRI from 3 stations. Gray bars show overlap.

3 Geometrical Distortion in MRI

Essential for MRI is to know the resonance frequency at each position within the FOV, to relate the frequency spectrum of the received RF impulses to the right spatial location. The Larmor equation relates precession frequency of nuclear spins ω to the magnetic field, composed of the static field B_0 and the slice and row selection gradient fields G_z

and G_x , respectively

$$\varpi(x, z) = \gamma(B_0 + x \cdot G_x + z \cdot G_z) \quad (1)$$

with γ the gyromagnetic ratio. Inhomogeneity of the static field or nonlinearity of gradient fields, more prominent farther away from the magnetic iso-center, lead to uncertainties, causing geometrical distortion artefacts. More specifically, these artefacts are referred to be *system-specific* artefacts in contrast to *patient-induced* artefacts arising from susceptibility effects, chemical shift, and flow [4]. Scanner with higher field strength and shorter bore magnets, the current trend in MRI, are more sensitive to distortion effects, putting its correction back into focus [4].

In the literature, mainly system-specific artefacts are discussed. Chang and Fitzpatrick [5] correct for B_0 distortion by acquiring two almost identical images only differing in the polarity of the read-out gradient. This enables for an exact correction of B_0 distortion. In follow-up articles Kannengiesser *et al.* [6] and Reinsberg *et al.* [7] refine this method by using deformable image registration techniques to make it also applicable to real MR images and not only phantom scans. This approach is hardly applicable in the clinical workflow because acquisition time would double.

Doran *et al.* [4] analyze the distortion with phantom scans and apply the deduced correction field to medical data. It is doubtful if corrections based on phantom scans make sense because it is not feasible to build a phantom duplicating a biological system, and distortions calculated at fiducial locations have to be interpolated to create a dense mapping, which limits the accuracy [5]. Our approach has the advantage that no additional images have to be acquired, enabling a seamless integration in the clinical workflow. Additionally, the correction is not limited to system-induced distortions but also covers patient-induced ones.

4 Deformable Mosaicing

In order to introduce our approach of deformable mosaicing, we define the two volumes to be stitched as $I_1 : \Omega_1 \subset \mathbb{R}^3 \rightarrow \mathbb{R}$ and $I_2 : \Omega_2 \subset \mathbb{R}^3 \rightarrow \mathbb{R}$. The overlapping domain is denoted as $\Omega_o = \Omega_1 \cap \Omega_2$. Since the overlap Ω_o is the only part where the two images share any information, a naive approach for the stitching could be defined as a minimization problem with respect to a certain distance/similarity measure $\rho(\cdot)$, or

$$\hat{\mathcal{T}}_{1,2} = \arg \min_{\mathcal{T}_{1,2}} \int_{\Omega_o} \rho(I_1(\mathcal{T}_1(\mathbf{x})) - I_2(\mathcal{T}_2(\mathbf{x}))) d\mathbf{x} \quad (2)$$

where $\mathbf{x} = (x, y, z)$ denotes a voxel position, and $\mathcal{T}_{1,2}$ are the parameters of the transformations \mathcal{T}_1 and \mathcal{T}_2 relating the two volumes in the spatial domain. The most common approach in pairwise registration is to assume that one of the two transformations is equal to the identity transformation. In our case, such an approach would lead to several problems: (i) through the selection of a moving and a fixed image, we would introduce a certain bias on the stitching result, (ii) since both volumes are distorted due to the inhomogeneous magnetic field in the overlap volume, none of them is actually representing a good reference for the stitching, and (iii) a registration performed only within the overlap may result in discontinuities with respect to the rest of the volumes.



Fig. 2: Synthetic example of a deformable stitching. The first and second image are to be stitched where both are significantly distorted. The initialization of our linear weighted average is shown in the third image. The horizontal gray lines indicate the borders of the overlap area. Fourth to sixth image is an illustration of the registration progress and the iterative improvement of the linear weighted average.

In order to overcome these problems, we propose an iterative simultaneous registration using a *linear weighted average*. The idea of the weighted average is to account for the underlying physical properties of increasing distortions towards the volume boundaries. Assuming that the boundary information is less reliable, we would like to reduce its influence to the registration.

4.1 Simultaneous Registration to Linear Weighted Average

Let us define another volume $S : \Omega_s$ on the union of the two volume domains $\Omega_s = \Omega_1 \cup \Omega_2$. The intensities of S are set using our average model, or

$$S(\mathbf{x}) = \begin{cases} f(\mathbf{x}), & \text{if } \mathbf{x} \in \Omega_o \\ I_1(\mathcal{T}_1(\mathbf{x})), & \text{if } \mathbf{x} \in \Omega_1 \setminus \Omega_2 \\ I_2(\mathcal{T}_2(\mathbf{x})), & \text{if } \mathbf{x} \in \Omega_2 \setminus \Omega_1 \end{cases} \quad (3)$$

where $f(\cdot)$ is a function computing the linear weighting in the overlap volume, or

$$f(\mathbf{x}) = (1 - h(\mathbf{x})) \cdot I_1(\mathcal{T}_1(\mathbf{x})) + h(\mathbf{x}) \cdot I_2(\mathcal{T}_2(\mathbf{x})). \quad (4)$$

The linear function $h(\cdot)$ has a range of $(0, 1)$ and is defined for the overlap domain Ω_o with respect to the stitching direction. In our application, this direction is usually along the head-feet axis which corresponds to the y-axis of our common 3D coordinate system for all the MRI volumes.

The setup for the deformable stitching and the initialization of the linear weighted average is illustrated in Fig. 2. We can reformulate the naive registration in Eq. 2 in order to pose a simultaneous registration based on the linear weighted average S . In terms of an energy function (which is to be minimized), we define

$$E_{\text{data}}(\mathcal{T}_{1,2}) = \sum_{i=1}^2 \int_{\Omega_o} \rho(S(\mathbf{x}) - I_i(\mathcal{T}_i(\mathbf{x}))) d\mathbf{x}. \quad (5)$$

In order to reduce the dimensionality of the problem, we consider Free Form Deformations [8] as the transformation model for the two images. A deformation grid $G : [1, K] \times [1, L] \times [1, M]$ is superimposed onto the volume domain Ω_s . By deforming the grid (with a 3D displacement vector \mathbf{d}_p for each control point) the underlying structures are aligned. The transformation of a voxel \mathbf{x} can be expressed using a combination

of basis functions, or

$$\mathcal{T}(\mathbf{x}) = \mathbf{x} + \mathcal{D}(\mathbf{x}) \quad \text{with} \quad \mathcal{D}(\mathbf{x}) = \sum_{\mathbf{p} \in G} \eta(|\mathbf{x} - \mathbf{p}|) \mathbf{d}_{\mathbf{p}} \quad (6)$$

where $\eta(\cdot)$ is the weighting function (based on cubic B-Splines) measuring the contribution of the control point \mathbf{p} to the displacement field \mathcal{D} .

Now, we can rewrite the objective function defined in Eq. 5 based on the two deformation grids G_1 and G_2 , or

$$E_{\text{data}}(\mathcal{T}_{1,2}) = \sum_{i=1}^2 \frac{1}{|G_i|} \sum_{\mathbf{p} \in G_i} \int_{\Omega_o} \hat{\eta}(|\mathbf{x} - \mathbf{p}|) \cdot \rho(S(\mathbf{x}) - I_i(\mathcal{T}_i(\mathbf{x}))) d\mathbf{x}. \quad (7)$$

where $\hat{\eta}(\cdot)$ computes the influence of a voxel \mathbf{x} to a control point \mathbf{p} . Such a function acts as a *projection* of the distance/similarity measure computed from the volume domain back to the coarser level of control points. Different definitions of the $\hat{\eta}(\cdot)$ have to be considered with respect to the used similarity measure. We use the normalized cross correlation (NCC) which is robust to intensity variations common in MRI. For statistical measures such as NCC, we define

$$\hat{\eta}(|\mathbf{x} - \mathbf{p}|) = \begin{cases} 1, & \text{if } \eta(|\mathbf{x} - \mathbf{p}|) > 0 \\ 0 & \text{otherwise} \end{cases}. \quad (8)$$

Basically, this function masks voxels influenced by a control point \mathbf{p} resulting in a local image patch centered at the control point. From this patch, a *local* similarity measure can then be computed. Definitions for voxel-wise measures such as sum of squared differences can be found in [9].

The simultaneous registration to an average should overcome the problems for the reference selection, mentioned before. This is very similar to atlas construction approaches where the average is used as the reference image in order to achieve an unbiased coordinate frame (e.g. for shape models) [10]. In addition, we try to account for the increasing distortions using a linear weighted average.

4.2 Optimization through Discrete Labeling

We propose to define the simultaneous registration as a discrete optimization problem. Discrete optimization has been recently shown to provide very good results in the case of standard pairwise registration [9]. Based on the previous assumptions, we define a set of discrete variables $G_{\text{mrf}} = G_1 \cup G_2$. Thus, each variable corresponds to a control point of one of the two deformation grids. Similar to [9], we consider a discrete set of labels $L = \{l^1, \dots, l^I\}$ corresponding to a quantized version of the deformation space $\Theta = \{\mathbf{d}^1, \dots, \mathbf{d}^I\}$. A label assignment $l_{\mathbf{p}}$ to a grid node \mathbf{p} is associated with displacing the node by the corresponding vector $\mathbf{d}^{l_{\mathbf{p}}}$. If a label is assigned to every node we get a *discrete labeling* \mathbf{l} . A popular model for representing discrete labeling problems are Markov Random Fields (MRFs) [11]. The general form of a first-order MRF is

$$E_{\text{mrf}}(\mathbf{l}) = \sum_{\mathbf{p} \in G_{\text{mrf}}} V_{\mathbf{p}}(l_{\mathbf{p}}) + \sum_{(\mathbf{p}, \mathbf{q}) \in \mathcal{E}_{\text{mrf}}} V_{\mathbf{pq}}(l_{\mathbf{p}}, l_{\mathbf{q}}) \quad (9)$$

where $V_{\mathbf{p}}(\cdot)$ are the unary potentials representing the data term, $V_{\mathbf{pq}}(\cdot, \cdot)$ are the pairwise potentials representing the smoothness term, and \mathcal{E}_{mrf} represents the neighborhood system represented by edges between nodes. We define the unary potentials (in iteration t) according to our data term $\forall \mathbf{p} \in G_i$ as

$$V_{\mathbf{p}}(l_{\mathbf{p}}) = \int_{\Omega_o} \hat{\eta}(|\mathbf{x} - \mathbf{p}|) \cdot \rho(S(\mathbf{x}) - I_i(\mathcal{T}_i^{t-1}(\mathbf{x}) + \mathbf{d}^{l_{\mathbf{p}}})) d\mathbf{x}. \quad (10)$$

The pairwise potentials encode a penalty term for assigning different labels to neighboring nodes. The FFD transformation model already inherits some implicit smoothness properties. Additionally, one can consider explicit regularizing constraints on the grid domain using the pairwise potentials. These are defined $\forall (\mathbf{p}, \mathbf{q}) \in \mathcal{E}_{\text{mrf}} \wedge \mathbf{p}, \mathbf{q} \in G_i$ as

$$V_{\mathbf{pq}}(l_{\mathbf{p}}, l_{\mathbf{q}}) = \lambda \left| (\mathcal{R}_i(\mathbf{p}) + \mathbf{d}^{l_{\mathbf{p}}}) - (\mathcal{R}_i(\mathbf{q}) + \mathbf{d}^{l_{\mathbf{q}}}) \right| \quad (11)$$

where λ denotes a weighting factor for the smoothness term and $\mathcal{R}_i(\cdot)$ back-projects the accumulated displacement field (of iteration $t - 1$) on the control point level, or

$$\mathcal{R}_i(\mathbf{p}) = \int_{\Omega_s} \hat{\eta}(|\mathbf{x} - \mathbf{p}|) \mathcal{D}_i^{t-1}(\mathbf{x}) d\mathbf{x}. \quad (12)$$

In contrast to the data term energy, the smoothness energy affects the whole stitching domain Ω_s . Such an approach together with the use of smooth FFD transformations guarantees continuous and seamless transitions between the overlapping and non-overlapping areas of the stitched volume while the actual stitch is (softly) constrained to the volume of overlap.

Many optimization algorithms exist for efficiently solving discrete labeling problems in forms of an MRF [11]. We use a recently proposed method called Fast-PD [12] which is also used in [9]. Due to the limited space, we refer the reader to the given references.

4.3 Iterative Multi-Scale Approach

We propose to embed the simultaneous registration into a common iterative multi-scale approach. The simultaneous registration of the two volumes I_1 and I_2 to the linear weighted average S is performed in a pyramidal setup where different levels of resolution for the volume as well as for the deformation grids are considered. On each level, several discrete labelings are computed where the set of displacements Θ is successively refined each time and the displacement fields are incrementally updated. After the registration converges, the linear weighted average S is recomputed and a new registration cycle is started. In Fig. 2 three of such cycles are illustrated for the case of synthetic data. Usually, only a few cycles are needed until the average shows no dramatic changes anymore.

5 Experimental Validation

We evaluate our method on 8 whole-body T1- and T2-weighted data sets from three different Siemens MR scanners: Avanto 1.5T, Trio 3T, and Espree 1.5T. The overlaps




Fig. 3: The two rows are magnifications of the stitching areas of the WB-MRI shown in Fig. 1. Left: Initial average. Middle: Final stitching result after 3 optimization cycles. Right: Reference scan where the overlap volume is centered within the MR scanner. Our method is able to reproduce similar smooth and continuous transitions as present in the reference images.

vary between 5 and 27 cm. For all stitching results, we obtained very positive feedback from our clinical experts in the radiology department who inspected the images visually. An example mosaic is shown in Fig. 1, consisting of three volumes having a FOV of $50 \times 50 \times 28 \text{ cm}^3$, a resolution of $448 \times 448 \times 35$ voxel, and an overlap of 5 cm. In Fig. 3, one can clearly see the influence of the distortion by regarding the initial average, and the improvement after the deformation by comparing the result to the reference scan. The final resolution for the 3 station stitch shown in Fig. 1 is $448 \times 1256 \times 35$ where the two stitches take together approximately 25 min of computational time.

To illustrate that the proposed method also works for varying overlaps, we show the stitching of 3 volumes for whole-spine MR, see Fig. 4. The first overlap is with 15.2 cm very large and our method arrives at producing a sharper average. The second one, with only 1.4 cm, shows discontinuities in the initial average, which are removed after deformable mosaicing.

6 Conclusion

Speeding up the acquisition for WB-MRI with large FOV images leads to significant distortions towards the boundaries. Methods for distortion correction proposed in the literature are not applicable to the WB imaging setup because they either elongate the workflow or only correct for specific system-induced distortions. So far, the overlap in WB-MRI was not used to correct for distortion. We propose the usage of simultaneous deformable registration in a mosaicing scenario, which has not yet been done before. Key for the simultaneous registration is the creation of a linear weighted average, each of the two images is registered to. Our experiments on synthetic and in-vivo data show the ability of the method to correct for distortions. The unaltered clinical workflow makes our approach very interesting for being integrated into further MR scanner generations.




Fig. 4: Stitching of 3 spine volumes. Top: initial average. Bottom: result. Gray bars indicate overlap.

References

1. Schmidt, G.P., Reiser, M.F., Baur-Melny, A.: Whole-body imaging of the musculoskeletal system: the value of MR imaging. *Skel Rad* **36**(12) (2007) 1109–1119
2. Lauenstein, T.C., Goehde, S.C., Herborn, C.U., Goyen, M., Oberhoff, C., Debatin, J.F., Ruehm, S.G., Barkhausen, J.: Whole-Body MR Imaging: Evaluation of Patients for Metastases. *Radiology* **233**(1) (2004) 139–148
3. Goyen, M.: *Real Whole Body MRI: Requirements, Indications, Perspectives*. McGraw-Hill (2007)
4. Doran, S., Charles-Edwards, L., Reinsberg, S., Leach, M.: A complete distortion correction for MR images: I. Gradient warp correction. *Physics in Medicine and Biology* **50**(7) (2005) 1343–1361
5. Chang, H., Fitzpatrick, J.: A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE TMI* **11**(3) (1992) 319–329
6. Kannengiesser, S., Wang, Y., Haacke, E.: Geometric distortion correction in gradient-echo imaging by use of dynamic time warping. *Magnetic Resonance in Medicine* **42**(3) (1999) 585–590
7. Reinsberg, S., Doran, S., Charles-Edwards, E., Leach, M.: A complete distortion correction for MR images: II. Rectification of static-field inhomogeneities by similarity-based profile mapping. *Phys Med Biol* **50**(11) (2005) 2651–2661
8. Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D.: Nonrigid registration using free-form deformations: application to breast mr images. *Medical Imaging, IEEE Transactions on* **18**(8) (1999) 712–721
9. Glocker, B., Komodakis, N., Paragios, N., Tziritas, G., Navab, N.: Inter and intra-modal deformable registration: Continuous deformations meet efficient optimal linear programming. In: *IPMI, Kerkrade, Netherlands* (July 2007)
10. Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* **23** (2004) 151–160
11. Li, S.Z.: *Markov random field modeling in image analysis*. Springer-Verlag New York, Inc. (2001)
12. Komodakis, N., Tziritas, G., Paragios, N.: Fast, approximately optimal solutions for single and dynamic mrfs. In: *CVPR*. (2007)