Correction of Fat-Water Swaps in Dixon MRI

Ben Glocker[1], Ender Konukoglu[2], Ioannis Lavdas[3], Juan Eugenio Iglesias[4],
Eric O. Aboagye[3], Andrea G. Rockall[3], Daniel Rueckert[1]

[1] Biomedical Image Analysis Group, Imperial College London, UK
[2] Computer Vision Laboratory, ETH Zurich, Switzerland
[3] Comprehensive Cancer Imaging Centre, Imperial College London, UK
[4] Translational Imaging Group, UCL, UK, and BCBL, Spain

Abstract. The Dixon method is a popular and widely used technique
for fat-water separation in magnetic resonance imaging, and today, nearly
all scanner manufacturers are offering a Dixon-type pulse sequence that
produces scans with four types of images: in-phase, out-of-phase, fat-
only, and water-only. A natural ambiguity due to phase wrapping and
local minima in the optimization problem cause a frequent artifact of
fat-water inversion where fat- and water-only voxel values are swapped.
This artifact affects up to 10% of routinely acquired Dixon images, and
thus, has severe impact on subsequent analysis. We propose a simple yet
very effective method, *Dixon-Fix*, for correcting fat-water swaps. Our
method is based on regressing fat- and water-only images from in- and
out-of-phase images by learning the conditional distribution of image
appearance. The predicted images define the unary potentials in a glob-
ally optimal maximum-a-posteriori estimation of the swap labeling with
spatial consistency. We demonstrate the effectiveness of our approach on
whole-body MRI with various types of fat-water swaps.

1 Introduction

Reliable fat-water separation in magnetic resonance imaging (MRI) has a wide
range of clinical applications [1]. As the fat signal has a relatively short T1 relax-
ation time, its bright appearance, without separation or suppression, can obscure
pathological structures and patterns of edema, inflammation, or enhancing tu-
mors [1]. Additionally, if separated, the fat signal can highlight pathologies such
as fatty tumors. The quantification of the amount of visceral adipose tissue is
another example requiring reliable separation of water from fat. There are sev-
eral techniques for fat suppression and fat-water imaging in MRI, and among
these chemical shift based fat-water separation is one of the most popular and
widely used ones [2]. This is also referred to as the Dixon method [3], originally
developed as a simple spectroscopic imaging technique using two images I and
O with a modified spin echo pulse sequence from which fat- and water-only im-
ages F and W are derived. The images I and O are acquired at time points
where the fat and water signal are *in-phase*, i.e. $I = W + F$, and *out-of-phase*,
i.e. $O = W - F$. Theoretically, fat- and water-only signals can be recovered by
$F = \frac{1}{2}(I - O)$ and $W = \frac{1}{2}(I + O)$. The original approach is sensitive to B_0

inhomogeneities, and hence, many improvements had been suggested, including three- and multi-point Dixon methods [4] which overcome the limitation of an inherent ambiguity in the fat and water signal. A breakthrough has been achieved by employing mathematical optimization for iterative decomposition, known as IDEAL [5]. Nowadays, all major scanner manufacturers offer one or multiple variants of a Dixon-based sequence for routine, clinical fat-water imaging.

Still, inherent to all Dixon methods is an artifact that causes fat-water swaps due to a natural ambiguity in the phase encoding and convergence of the employed optimization method to local minima. Even with sophisticated unwrapping algorithms this artifact cannot be completely avoided [1]. If undetected, those swaps can severely impact subsequent processing and analysis steps, for example, yielding wrong estimates of body fat. In a recent study, the impact of fat-water inversion has been investigated in the context of PET/MR, where Dixon images are used for attenuation correction [6]. The study found that 8% (23 of 283) of the images were affected by fat-water swaps. The average fraction of fat in the brain calculated from images with no fat-water swap was 13%, and 56% in the scans with inversions demonstrating the severe impact that this artifact can have on analysis. In our own database consisting of 46 subjects with whole-body Dixon MRI we found 10 cases exhibiting various types of fat-water swaps. Visual examples are shown in Fig. 1. While in smaller studies, visual quality control can be established to detect (and possibly discard) affected scans, this is not possible in large-scale studies with many hundreds or even thousands of subjects. Automated methods for detecting whether fat-water swaps are present are required. Furthermore, due to the high costs for imaging it is desirable to be able to include as many subjects as possible for further analysis, and thus, a method that can not only detect but correct fat-water swaps is needed.

1.1   Related work

Robust estimation of the inherent fat-water separation is an active field of research, and an overview is provided in [7]. There are several extensions [8] and alternatives to the original IDEAL method, including graph-cut formulations [7] and other discrete optimization approaches [9,10]. All require access to the original echo sequences or a B0 field map [11] as part of their optimization procedure. However, these data are not available after the scanning has been completed. In fact, most Dixon protocols yield exactly four types of images, irrespective of the underlying sequence, optimization procedure and number of echo times. The in-phase, out-of-phase, fat- and water-only images are what is accessible once the data have been reconstructed and stored in the imaging archive. To the best of our knowledge, there exist no technique that allows correction of fat-water swaps retrospectively based on just those derived images. Such a technique, however, would be very valuable as it can be applied to already acquired datasets, in particular in clinical studies, where the presence of fat-water swaps is often only detected during subsequent analysis. Being able to correct such corrupted data is important in clinical research as data are expensive to acquire and samples sizes are calculated tightly to reduce scanning time and costs.
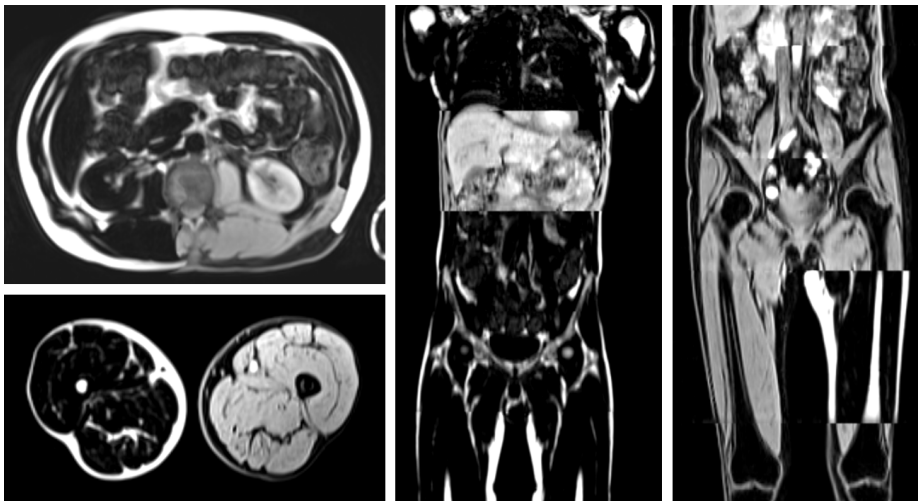
Fig. 1: Examples of different types of fat-water swaps in Dixon MRI. Swaps can affect local regions as seen in on the axial slices on the left for kidneys and legs. In whole-body MRI, entire stations can be affected as seen in the middle image, or partially affected, as seen for the lower extremities on the right.

1.2 Contribution

We propose a simple yet very effective approach for correcting fat-water swaps that can be applied retrospectively on already acquired data taking only the four derived images, in-phase, out-of-phase, fat- and water-only, as input. Our method makes use of image synthesis, i.e., a regression problem dealing with the prediction of images of a certain modalities, given input image(s) of other modalities. In our case, we learn the conditional distribution of image appearance of the fat- and water-only images given the in- and out-of-phase images. We utilize a training database of fat-water swap-free Dixon images. The predicted fat- and water-only images yield a noisy, but swap-free estimate that is employed in a voxel-wise energy term within a swap-labeling problem that is solved via graph-cuts. The details of our method are given in the following.

2 Dixon-Fix: Correction of Fat-Water Swaps

The core of our method for detecting and correcting fat-water swaps in Dixon MRI is based on an image synthesis component for predicting fat- and water-only images, F and W, from the in-phase and out-of-phase images, I and O. Given a training database $T = \{\mathcal{D}_i\}_{i=1}^m$ of m Dixon MRIs $\mathcal{D}_i = \{I_i, O_i, F_i, W_i\}$ where (F_i, W_i) are known to be swap-free, we can learn the conditional distribution $p(F, W | I, O)$. For a set of test images $(I_{\text{test}}, O_{\text{test}})$, we can then predict a potentially noisy, but swap-free estimation $(F_{\text{test}}^\star, W_{\text{test}}^\star)$ by inferring the most likely
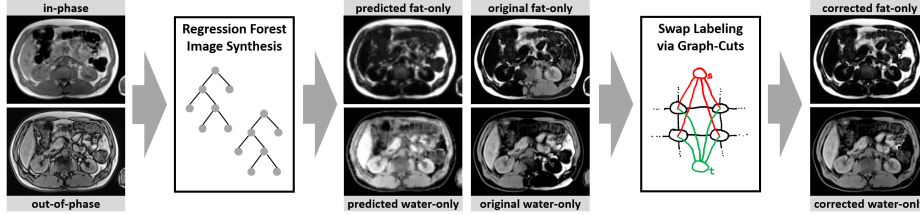
Fig. 2: Overview of the Dixon-Fix pipeline with the two main components of image synthesis to predict swap-free fat- and water-only images, and the swap labeling solved with graph-cuts that yields the corrected images.

images according to $(F^{\star}_{\text{test}}, W^{\star}_{\text{test}}) = \arg\max_{(F,W)} p(F, W | I_{\text{test}}, O_{\text{test}})$. For learning this conditional distribution, we employ Regression Forests which have been shown to yield good synthesis results in the context of diffusion MRI [12,13]. Using the predicted fat- and water-only images, we then formulate a binary labeling task in terms of an energy minimization problem of a pairwise Markov random field (MRF). Both, the image synthesis based on Regression Forests and the MRF labeling formulation are detailed in the following. Figure 2 provides an overview of the whole processing pipeline for correcting fat-water swaps.

2.1   Prediction of Swap-Free Fat- and Water-only Images

Regression Forests are an attractive choice for image synthesis as they can efficiently handle the multivariate, high-dimensional input data and yield a probabilistic, joint estimate of the fat- and water-only voxel values. Additionally, it is straightforward to incorporate both local and contextual information based on the popular, randomized, offset box-features which are efficiently implemented using integral images. We follow the basic Regression Forest methodology as described in [14]. We employ a leaf node predictor model based on multivariate Gaussians which summarize the statistics over the incoming training data. This model then yields predictions for every voxel v of the fat- and water-only values corresponding to the means $F^{\star}_{\mu}(v)$ and $W^{\star}_{\mu}(v)$, together with an estimate of the uncertainty corresponding to the standard deviations $F^{\star}_{\sigma}(v)$ and $W^{\star}_{\sigma}(v)$. The greedy learning procedure during which the randomized decision trees are constructed is driven by an objective function that aims to minimize the variance in each internal split node. Those split nodes are associated with binary tests applied on the most discriminate feature responses. At test time, the predictions from individual trees are aggregated by averaging.

2.2   Estimation of Swap-Labeling using Graph-Cuts

The image synthesis yields a swap-free estimate for the fat- and water-only images, denoted as $(F^{\star}_{\text{test}}, W^{\star}_{\text{test}})$. To determine which voxels in the original images $(F_{\text{test}}, W_{\text{test}})$ need to be swapped, we formulate a binary labeling problem as a

Markov random field where each voxel v is associated with a binary random variable x_v that takes values $\{0, 1\}$, with *zero* meaning 'no-swap', and *one* meaning 'swap'. The probability of a voxel being swapped $p(x_v|F(v), W(v), F^\star(v), W^\star(v))$ given the original and predicted fat- and water-only voxel values $F(v)$, $W(v)$, $F^\star(v)$, $W^\star(v)$ is defined with respect to similarities between those values. We are omitting the $_\text{test}$ subscript to avoid clutter. Mathematically, this can be defined as an MRF unary potential function as

$$\psi_v(x_v) = \begin{cases} \frac{|F(v)-F_\mu^\star(v)|}{F_\sigma^\star(v)} + \frac{|W(v)-W_\mu^\star(v)|}{W_\sigma^\star(v)} & \text{if } x_v = 0 \\[2ex] \frac{|F(v)-W_\mu^\star(v)|}{W_\sigma^\star(v)} + \frac{|W(v)-F_\mu^\star(v)|}{F_\sigma^\star(v)} & \text{if } x_v = 1 \end{cases} . \tag{1}$$

The potential function is a robust version of an intensity difference term that takes the voxel-wise uncertainty, $F_\sigma^\star(v)$ and $W_\sigma^\star(v)$, from the image synthesis into account. The potential function determines the cost of assigning 'non-swap' ($x_v = 0$) and 'swap' ($x_v = 1$) labels depending on whether the original voxel values are closer to the predicted fat- or water-only voxel values. To encourage spatial consistency, we add a simple pairwise potential function based on the Potts model which is defined for neighboring voxels (v, w) in a 6-neighborhood as $\psi_{vw}(x_v, x_w) = \lambda \cdot \mathbb{1}_{x_v \neq x_w}$. The pairwise function returns a penalty λ if the label assignments to voxels v and w are different, i.e. $x_v \neq x_w$, and it returns a zero-penalty otherwise. The total MRF energy for a swap-labeling $\mathbf{x}$ is then defined as the sum of the unary and pairwise potentials, $E(\mathbf{x}) = \sum_v \psi_v(x_v) + \sum_{(v,w)} \psi_{vw}(x_v, x_w)$. The globally optimal solution in terms of the maximum-a-posteriori estimate $\hat{\mathbf{x}} = \arg\max_\mathbf{x} E(\mathbf{x})$ is efficiently computed using graph-cuts.

3   Experiments

To evaluate the effectiveness of our correction method, we utilize a database of whole-body Dixon MRIs from 46 subjects. The data have been visually inspected and 10 cases have been identified with fat-water swaps. We normalize for intensity variation by matching intensity means to a fixed value. From the 36 swap-free cases, we randomly select 23 for training the Regression Forest with 50 trees, maximum depth 30, and a minimum of 4 samples per leaf as stopping criterion. A large pool of box-feature parameters ($n = 10,000$) is randomly generated. Parameters include the image channels, allowing cross-channel box difference features, and the side lengths and offsets of the 3D cuboids. We evaluate 200 randomly selected features on-the-fly for all incoming training samples at each split node. Ten equidistant thresholds in the interval defined by the minimum and maximum feature response are tested, and the overall best feature/threshold pair yielding largest decrease of the trace of the covariance matrix over the two-dimensional target variable is selected. This greedy, randomized optimization has shown to yield decision trees with good generalization ability.

Given a multi-channel input with in- and out-of-phase images, the forest predicts jointly the appearance of swap-free fat- and water-only images. The

swap labeling is determined efficiently through multi-resolution graph-cuts. A dense 3D MRF is defined first on low resolution images with 6mm voxel spacing. The resulting labeling defines the region for a full resolution non-regular, sparse MRF with nodes corresponding to voxels labeled as 'swap' in the first stage. The result of the second stage is an accurate swap labeling map used to generate the corrected fat- and water-only images. In all experiments, we fix the regularization parameter $\lambda = 5$, empirically chosen, and its optimal value depends on the range of MR intensities. The average running time of the synthesis is about 10 minutes for a whole-body scan, and the multi-resolution graph-cut takes less than 5 seconds. The fat-water swaps in all 10 cases are successfully corrected. Visual examples with original, corrupted images, swap labeling maps, and corrected images are shown in Fig. 3. The examples illustrate the variety of swaps that can be handled, including one case with multiple swaps across the whole-body image. When tested on the 13 swap-free images, we find that in all cases some voxels are incorrectly predicted as 'swap'. Thus, with the current approach it is not possible to confirm whether an image is swap-free. The incorrectly labeled voxels, however, either correspond to background regions or form small patches in noisy, unstructured and less important areas which both appear dark in the fat- and water-only images and therefore cause confusion in the MRF unaries.

4   Conclusion

The Dixon-Fix method provides an effective way for correcting fat-water swaps which is demonstrated qualitatively on whole-body MRI. Quantitative evaluation is difficult due to unknown ground truth. In future work synthesized swaps will be considered for this purpose, however, it should be noted that generating realistic swaps is difficult as the underlying process that causes these artifacts is quite complex. As our test cases exhibit a large variety of swaps in different areas, such as head/neck, upper and lower extremities, and internal organs, it is expected that the method works well in other applications, such as spine and brain. The two main components of our pipeline are based on established techniques, Regression Forests and graph-cuts. Others could be considered, in particular, for the synthesis, which is an active area of research with many recently proposed methods (see [15] for an overview). Probabilistic approaches might be considered in particular, as the predictions should come with confidence estimates for integration into the MRF energy. A direction for future work concerns the regularization term. The location-agnostic Potts model with a global weighting factor might not be optimal for all cases of fat-water swaps. A learned prior that takes contextual information and expected size and shape of the swap area into account could be explored. This might also provide effective means of removing the false positive swap predictions in swap-free images. To facilitate further development, the source code of our Dixon-Fix implementation is made publicly available on `http://biomedia.doc.ic.ac.uk/software/`.

Acknowledgement

References

1. Bley, T.A., Wieben, O., François, C.J., Brittain, J.H., Reeder, S.B.: Fat and water magnetic resonance imaging. J Magnetic Resonance Imaging **31**(1) (2010) 4–18
2. Ma, J.: Dixon techniques for water and fat imaging. J Magnetic Resonance Imaging **28**(3) (2008) 543–558
3. Dixon, W.T.: Simple proton spectroscopic imaging. Radiology **153**(1) (1984) 189–194
4. Glover, G.H.: Multipoint dixon technique for water and fat proton and susceptibility imaging. J Magnetic Resonance Imaging **1**(5) (1991) 521–530
5. Reeder, S.B., Pineda, A.R., Wen, Z., Shimakawa, A., Yu, H., Brittain, J.H., Gold, G.E., Beaulieu, C.H., Pelc, N.J.: Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): application with fast spin-echo imaging. MR in medicine **54**(3) (2005) 636–644
6. Ladefoged, C.N., Hansen, A.E., Keller, S.H., Holm, S., Law, I., Beyer, T., Højgaard, L., Kjær, A., Andersen, F.L.: Impact of incorrect tissue classification in dixon-based MR-AC: fat-water tissue inversion. EJNMMI Phys **1**(1) (2014) 101
7. Hernando, D., Kellman, P., Haldar, J., Liang, Z.P.: Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm. MR in medicine **63**(1) (2010) 79–90
8. Sharma, S.D., Artz, N.S., Hernando, D., Horng, D.E., Reeder, S.B.: Improving chemical shift encoded water–fat separation using object-based information of the magnetic field inhomogeneity. MR in medicine **73**(2) (2015) 597–604
9. Berglund, J., Ahlström, H., Johansson, L., Kullberg, J.: Two-point dixon method with flexible echo times. MR in medicine **65**(4) (2011) 994–1004
10. Berglund, J., Kullberg, J.: Three-dimensional water/fat separation and T2* estimation based on whole-image optimization: Application in breathhold liver imaging at 1.5T. MR in medicine **67**(6) (2012) 1684–1693
11. Narayan, S., Kalhan, S.C., Wilson, D.L.: Recovery of chemical estimates by field inhomogeneity neighborhood error detection (REFINED): Fat/water separation at 7 tesla. J Magnetic Resonance Imaging **37**(5) (2013) 1247–1253
12. Jog, A., Roy, S., Carass, A., Prince, J.L.: Magnetic resonance image synthesis through patch regression. In: Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on, IEEE (2013) 350–353
13. Alexander, D.C., Zikic, D., Zhang, J., Zhang, H., Criminisi, A.: Image quality transfer via random forest regression: applications in diffusion MRI. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014. Springer (2014) 225–232
14. Criminisi, A., Shotton, J., Konukoglu, E.: Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. Foundations and Trends® in Computer Graphics and Vision **7**(2–3) (2012) 81–227
15. Vemulapalli, R., Van Nguyen, H., Kevin Zhou, S.: Unsupervised cross-modal synthesis of subject-specific scans. In: Proceedings of the IEEE International Conference on Computer Vision. (2015) 630–638
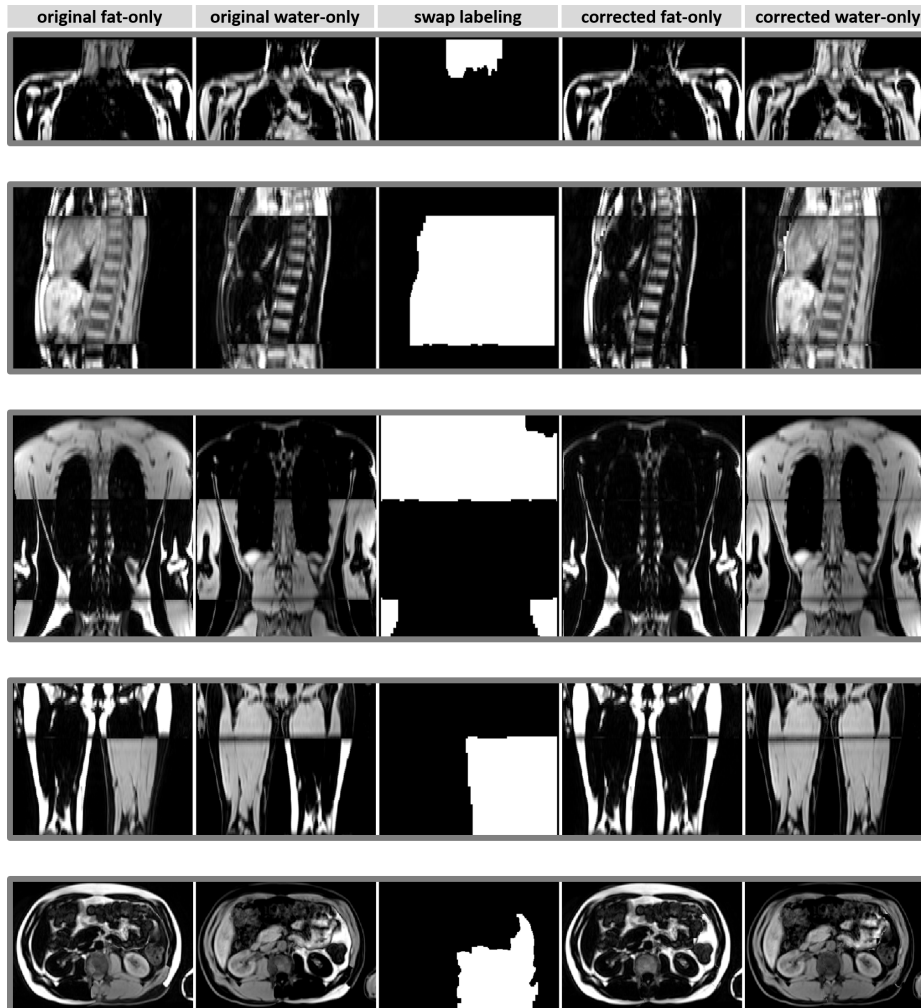
Fig. 3: Visual results of successful fat-water swap correction for five different cases. **First row** illustrates a case with a fat-water swap in the head/neck area. **Second row** shows a case of a whole station swap that is common in whole-body Dixon MRI. **Third row** shows a case with a mix of whole station swap and local swaps in the lower arm region. **Fourth row** shows successful correction of a left/right image swap, while the **last row** shows a challenging case of local, complex swaps of internal anatomical structures. Note that horizontal stripy artifacts are results of the stitching method used for generating whole-body images from individual stations, and not a result of the fat-water swap correction.