



# Correction of Susceptibility Distortion in EPI: A Semi-supervised Approach with Deep Learning

Antoine Legouhy<sup>1,2</sup>(✉), Mark Graham<sup>3</sup>, Michele Guerreri<sup>1,2</sup>, Whitney Stee<sup>4,5</sup>,  
Thomas Villemonteix<sup>4,5</sup>, Philippe Peigneux<sup>4,5</sup>, and Hui Zhang<sup>1</sup>

<sup>1</sup> Centre for Medical Image Computing, Department of Computer Science,  
University College London, London, UK  
a.legouhy@ucl.ac.uk

<sup>2</sup> AINOSTICS Ltd., Manchester, UK

<sup>3</sup> Department of Biomedical Engineering, School of Biomedical Engineering  
and Imaging Sciences, King's College London, London, UK

<sup>4</sup> UR2NF-Neuropsychology and Functional Neuroimaging Research Unit Affiliated  
at CRCN - Centre for Research in Cognition and Neuroscience,  
Université Libre de Bruxelles (ULB), Brussels, Belgium

<sup>5</sup> UNI-ULB Neuroscience Institute, Université Libre de Bruxelles (ULB),  
Brussels, Belgium

**Abstract.** Echo planar imaging (EPI) is the most common approach for acquiring diffusion and functional MRI data due to its high temporal resolution. However, this comes at the cost of higher sensitivity to susceptibility-induced  $B_0$  field inhomogeneities around air/tissue interfaces. This leads to severe geometric distortions along the phase encoding direction (PED). To correct this distortion, the standard approach involves an analogous acquisition using an opposite PED leading to images with inverted distortions and then non-linear image registration, with a transformation model constrained along the PED, to estimate the voxel-wise shift that undistorts the image pair and generates a distortion-free image. With conventional image registration approaches, this type of processing is computationally intensive. Recent advances in unsupervised deep learning-based approaches to image registration have been proposed to drastically reduce the computational cost of this task. However, they rely on maximizing an intensity-based similarity measure, known to be suboptimal surrogate measures of image alignment. To address this limitation, we propose a semi-supervised deep learning algorithm that directly leverages ground truth spatial transformations during training. Simulated and real data experiments demonstrate improvement to distortion field recovery compared to the unsupervised approach, improvement image similarity compared to supervised approach and precision similar to TOPUP but with much faster processing.

**Keywords:** Deep learning registration · Susceptibility distortion correction · Semi-supervised learning

## 1 Introduction

Echo planar imaging (EPI) is the most common approach for acquiring diffusion and functional MRI data due to its high temporal resolution which both reduces the influence of motion and allows the acquisition of a large number of volumes in a time frame amenable to neuroscientific and clinical research. This is, however, at the cost of higher sensitivity to susceptibility-induced  $B_0$  field inhomogeneities around interfaces of air, bone, and soft tissue. This leads to severe geometric distortions in the form of local expansions or contractions along the phase encoding direction (PED), breaking alignment with the corresponding anatomical scans and corrupting subsequent diffusion model fitting or tractography. Moreover, the effect of EPI susceptibility distortions radiate over the whole image, introducing systematic alterations even far from the apparent hot spots [6].

To tackle this problem, the strategy that has proved most effective is to acquire an extra scan with identical settings, except for an opposite PED [1] (also referred as blip up, blip down). It produces an analogous image with reverse distortion: expansions where there were contractions and vice versa. One can then apply non-linear image registration, with a transformation model constrained along the PED, to estimate the shift in voxel coordinates that undistorts the image pair and generates a distortion-free image. Standard implementations of this strategy, e.g. TOPUP [4] in FSL, are computationally intensive. They rely on traditional image registration methods that align each new pair of images with a separate iterative optimization. A comparison of such algorithms is available in [2]. It was proposed in [3] to synthesize a  $b = 0$  image through deep learning thus allowing to use TOPUP with the economy of the opposite PED acquisition, but with no gain in terms computational time.

Recently, image registration based on deep learning (DL) architectures, convolutional neural networks (CNN) in particular, have been developed. With the investment of an upfront cost during training, test images can be registered in one-shot almost instantaneously. In the same vein as the unsupervised Voxelmorph [16] framework for anatomical images, it has led to the development of fast EPI distortion correction with DL-powered registration. In [8] and [7], similarly to traditional registration, the optimization relies on maximizing an intensity similarity measure, known to be suboptimal surrogate measures of image alignment [9]. Also, [8] does not integrate intensity modulation to account for signal stretchings and pile-ups associated to geometric expansions and contractions. This aspect is given due consideration in [7], but the network is based on a 2D architecture likely to miss volumetric characteristics. In [5], more reliable fiber orientation distribution (FOD) features are used, although this is intended only to be used as secondary correction of the distortion residuals after an external primary one.

To address those limitations, we propose a 3D semi-supervised DL algorithm that directly leverages ground truth spatial transformations during training. We hypothesise that constraining DL model training with the most direct representations of spatial correspondence will significantly improve the fidelity of the recovered spatial correspondence during testing. Also we integrate Jacobian intensity modulation when constructing the undistorted images.

We performed two experiments to evaluate the proposed method against the unsupervised approach similar to [8], but also against a fully transformation supervised one. The first experiment makes use of real data to assess the practical performance. For this experiment, transformations produced by TOPUP are used as ground truth for training and testing. The second experiment makes use of simulated data, generated by DW-POSSUM [1], a realistic Spin-Echo EPI simulator. This sets up a scenario where we have genuine ground truth for both undistorted image and the distortion-inducing deformation, and allows the comparison with TOPUP.

The code used to implement the proposed semi-supervised model, as well as the unsupervised and field supervised ones (and more), is available in the open source *sudistoc*<sup>1</sup> repository.

## 2 Background

### 2.1 Distortion Model

Susceptibility-induced EPI geometric distortion is well understood [11]. For such multi-slice acquisitions, distortion due to  $B_0$  field inhomogeneities is negligible along the frequency encoding direction. Its effect can thus be parametrized by a unidirectional deformation field  $V$  shifting voxel coordinates  $x$  along the PED. For the typical PED from posterior to anterior, we have  $T_+(x) = x + V(x)$ . If the PED is reversed, an opposite displacement field will result:  $T_-(x) = x - V(x)$ . Henceforth, we will refer to this as the opposite symmetry constraint for the two displacement fields.

Without loss of generality, we will refer to  $T_+$  as the forward transform, and the corresponding distorted image as  $I_+$ . Likewise,  $T_-$  will be referred to as the backward transform and the corresponding distorted image as  $I_-$ . The latent undistorted image will be denoted as  $\hat{I}$ . Following [12], it can be expressed both in terms of the forward and backward images following:

$$\begin{cases} \hat{I}(x) \sim J_{T_+} \cdot I_+ \circ T_+(x) \\ \hat{I}(x) \sim J_{T_-} \cdot I_- \circ T_-(x) \end{cases} \quad (1)$$

where  $J_{T_+}$  (resp.  $J_{T_-}$ ) is the Jacobian determinant of  $T_+$  (resp.  $T_-$ ),  $\circ$  denoting composition and  $\cdot$  element-wise multiplication.  $J_T$  encodes the local expansion (if  $|J_T| \in [1, +\infty)$ ) or contraction (if  $|J_T| \in (0, 1]$ ) properties of the transformation and will modulate the resulting intensities accordingly [13].

### 2.2 Distortion Correction Using Image Registration

Under the above distortion model, it is evident image registration can be used to estimate the transformations for correcting the distorted image pair. Equation (1) suggests the correction can be formulated as the following image registration problem:

$$\arg \min_V C(J_{T_+} \cdot I_+ \circ T_+, J_{T_-} \cdot I_- \circ T_-) \quad (2)$$

<sup>1</sup> *sudistoc*: <https://github.com/CIG-UCL/sudistoc>.

where  $C$  is a dissimilarity criterion between the two corrected images (from  $I_+$  and  $I_-$  respectively). The mean squared error (MSE) between intensities is particularly well suited in this case since we are dealing with the same subject, the same modality, the same acquisition parameters (with the exception of PED); we therefore expect almost identical intensities at endpoint. The sought transformations ( $T_+$  and  $T_-$ ) are parametrized by  $V$  which, even though constrained to be unidirectional, can have a large number of degrees of freedom: up to the number of voxels of the image.

As noted in the introduction, this image registration problem is currently typically solved via computationally intensive iterative optimization, with each new image pair solved completely independently. Recent advances in DL-based image registration have recently been leveraged to substantially accelerate this task by replacing iterative optimization with a one-shot computation [14, 15]. The idea behind DL-based image registration is to learn a model that can predict from an image pair the transformation that put them into correspondence. During training, model parameters are tuned to output an optimal transformation, for each training sample, that maximises either some similarity between transformed images (unsupervised) or the resemblance to the corresponding ground truth transformation (supervised). While the training may be computationally expensive, once completed, new image pairs can be registered almost instantaneously. The most popular publicly available implementation is VoxelMorph [16, 17] which provides an unsupervised framework, making use of a U-Net convolutional neural network (CNN) architecture. This framework was recently exploited for EPI distortion correction [8]. However it has now been limited to optimisation over purely intensity-based losses and does not embed intensity modulation.

### 3 Method


We implement the proposed semi-supervised approach using the VoxelMorph framework. The framework must be adapted 1) to predict a spatial transform pair with opposite symmetry, 2) to constrain predicted spatial transforms to be unidirectional along the PED, 3) to support the image registration formulation represented by Eq. (2) that includes Jacobian intensity modulation, 4) to enable semi-supervision with ground truth spatial transforms, and 5) to handle weight maps that modulate the contribution of each voxels according to anatomical regions of interest. Points 3, 4 and 5 are not present in the work from [8]. The details of the model are described below.

#### 3.1 Model Architecture

The model is organised into three sequential blocks (Fig. 1). The first block takes a distorted image pair as input and outputs a unidirectional vector field which is used to produce a forward and a backward transformation. Those, together

with the associated input images, are fed to a resampling block that reconstruct undistorted images by interpolating from transformed coordinates.

The CNN block is a U-Net as in VoxelMorph, but its input and output are different. Instead of any arbitrary image pair, it must be given a pair of analogously distorted images  $\{I_+, I_-\}$  as input. Instead of a vector field, it outputs a scalar field  $V$  characterizing constrained displacements along a single direction (PED). As in VoxelMorph, all the trainable parameters of the model are contained in this block.



**Fig. 1.** Diagram of the proposed semi-supervised distortion correction network. Portions specific to unsupervised are highlighted in blue, whereas the ones associated to supervised are in red. (Color figure online)

To preserve the opposite symmetry constraint induced by the PED reversal, the pair of forward and backward transformations  $\{T_+, T_-\}$  are built from the estimated field  $V$  following:  $T_+ = \text{Id} + V$  and  $T_- = \text{Id} - V$ .

Resampling block implements image warping. Unlike VoxelMorph and [8], our implementation includes intensity modulation, which allows us to account for signal pile-up in the presence of contraction and signal reduction in the presence of expansion. It takes  $I_+$  (resp.  $I_-$ ) and  $T_+$  (resp.  $T_-$ ) as input to produce undistorted, intensity-modulated images  $J_{T_+} \cdot I_+ \circ T_+$  (resp.  $J_{T_-} \cdot I_- \circ T_-$ ). This requires computing the Jacobian determinant of the transformations.

### 3.2 Models

The model architecture described above is used to implement an unsupervised and a transformation supervised model, as baselines for comparison, and the proposed semi-supervised model.

A diagram illustrating the different models can be found in Fig. 1. The parts exclusive to the unsupervised model, comparable to [8], are highlighted in blue. It includes the resampling block that unwarpes the input images with the estimated field allowing their comparison (unsupervised loss). The red parts are exclusive to the supervised model. It includes the ground truth field that is compared to the estimated one (supervised loss). The black parts are common to both. It

includes the distortion field estimation (a regularization loss is computed on it). The semi-supervised model encompasses both the unsupervised and supervised components.

The unsupervised loss  $\mathcal{L}_U$  is an image similarity metric between the two undistorted images that have undergone Jacobian intensity modulation. As mentioned in Sect. 2.2 the MSE is well indicated and Eq. 1 leads to:

$$\mathcal{L}_U = \sum_{i=1}^n w_i (J_{T_+} \cdot I_+ \circ T_+(x_i) - J_{T_-} \cdot I_- \circ T_-(x_i))^2 \quad (3)$$

This loss is only required for the unsupervised and semi-supervised models.

The supervised loss is a distance between the estimated distortion field  $V$  and the ground truth one  $\hat{V}$ . The MSE is also well adapted when dealing with displacement vectors as it corresponds to an average of geometrical distances, it's a direct quantitative measure of the goodness of the registration:

$$\mathcal{L}_S = \sum_{i=1}^n w_i (V(x_i) - \hat{V}(x_i))^2 \quad (4)$$

This loss is only required for the supervised and the semi-supervised model.

A regularization loss  $\mathcal{L}_R$  is also use on the estimated distortion field  $V$  to encourage smoothness:

$$\mathcal{L}_R = \frac{1}{n} \sum_{i=1}^n \|\nabla_V(x_i)\|_F^2 \quad (5)$$

where  $\nabla_V$  is the Jacobian of the field  $V$  and  $\|\cdot\|_F$  is the Frobenius norm. This loss is computed for all models.

For each models, the overall loss is a sum of the ones above, weighted to account for large order-of-magnitude differences in different loss terms.

To prevent the learning process to be influenced by meaningless information from background (which represent the majority of the image!), another kind of weighting, spatially this time, occurs when computing the unsupervised and the supervised losses. It corresponds to the  $w_i$  in Eq. 3 and Eq. 4. The contribution of each voxels is modulated such as only areas of interest (typically just the brain) contribute to the loss, ignoring the background. This is not present in Voxelmorph and [8].

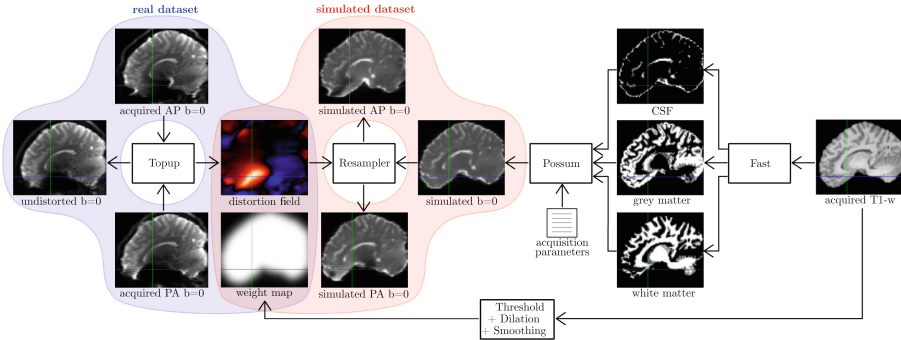
## 4 Evaluations

The idea is to evaluate how well, having processed a portion of a dataset with a regular tool (here TOPUP), one can use those to train a DL model that will rapidly process the rest of the dataset. Corrections from three DL registration approaches are engaged for comparison: unsupervised, transformation supervised and semi-supervised. We will perform experiments on two datasets: 1- A real dataset with TOPUP outputs as ground truths, 2- A synthetic dataset with absolute ground truths allowing to integrate TOPUP to the comparison.

## 4.1 Datasets

We acquired anatomical and diffusion weighted data as part of an ongoing study investigating memory learning and consolidation. We have, for 60 healthy subjects, T1-weighted images and distorted EPI  $b = 0$  image pairs from opposite PED acquisition that are antero-posterior (AP) and postero-anterior (PA).

Distorted EPI images have been processed through TOPUP to obtain undistorted images and associated distortion fields.



**Fig. 2.** Preprocessing steps involved in creating the real and the simulated datasets.

In addition, simulated data were produced using DW-POSSUM [1] (an extension of FSL POSSUM [10]), a realistic Spin-Echo EPI simulator. It takes as input tissue segmentations, MR parameters associated with these tissues and a pulse sequence, and produces an EPI image by solving the Bloch equations. We obtained the tissue segmentations by processing the anatomical scans with FSL FAST, producing probability maps for grey matter, white matter and cerebrospinal fluid. We then applied the distortion estimated by TOPUP on the real EPI data of each subject to their corresponding simulated undistorted images to create new synthetic pairs of AP and PA distorted images.

Acquired, TOPUP processed and simulated data have been used to produce two datasets, denoted as real and simulated, in order to cover various experimental configurations. A diagrammatic representation of the processing paths followed to obtain the two datasets is presented in Fig 2.

- In the real dataset: the inputs of the model are the acquired AP and PA images, the ground truth distortion fields used at training and for evaluation are the ones from TOPUP, the ground truth images used only for evaluation are the undistorted images from TOPUP. The advantage of this dataset is that it is made of real data and all the artifacts it implies. The drawback is that TOPUP is used as ground truth and therefore cannot be included in the comparison.

- In the simulated dataset: the inputs of the model are the simulated AP and PA images, the ground truth distortion fields used at training and for evaluation are the ones from TOPUP (same as for the real dataset), the ground truth images used only for evaluation are simulated non-distorted images. The advantage of this dataset is the synthetic, absolute nature of the ground truths allowing any algorithm comparison. The drawback is that the simulation process is not able to reproduce all the artifacts induced by a real acquisition.

For both, the weight maps have been computed by thresholding (binary brain mask), then dilating (3 voxels), then smoothing (Gaussian,  $\sigma = 6$  mm) the T1-weighted images.

We divided the same way the two datasets into training ( $n = 40$ ), validation ( $n = 10$ ) and testing ( $n = 10$ ) samples. Although it was shown in [16] that Voxelmorph-like architectures can achieve decent registration with only ten or so training subjects.

## 4.2 Models

We compared 3 deep learning approaches: unsupervised, transformation supervised and semi-supervised. Each model is trained and assessed separately for real and simulated experiments. To form the overall loss, the unsupervised loss was attributed a weight 200 000 (unsupervised and semi-supervised models only), the supervised loss was attributed a weight 300 (supervised and semi-supervised models only) and the regularization loss was attributed a weight 1 (all models). Each model was trained for 1000 epochs, with the epoch giving the best validation loss kept. Learning rate was set to  $10^{-4}$ . The training time for each model was about 3 h on a GPU.

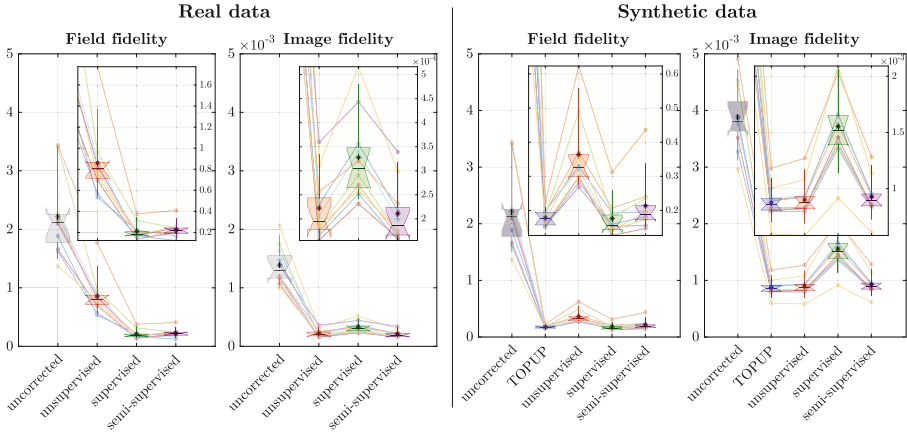
## 4.3 Assessment Metrics

For both real and simulated data experiments, their respective unseen test data were used to quantitatively assess the performance of the proposed semi-supervised approach against the current unsupervised approach. The assessment made use of the following metrics: 1) Image fidelity: MSE between the estimated undistorted versions of the AP and PA images and the corresponding ground truth undistorted image. 2) Field fidelity: MSE between the estimated distortion field and the corresponding ground truth distortion field (expressed in fraction of the voxel size which is 2 mm). These mean measures were weighted by the weight maps to ignore the background.

## 5 Results

Table 1 and Fig. 3 summarise the evaluation results for both real and simulated data experiments.



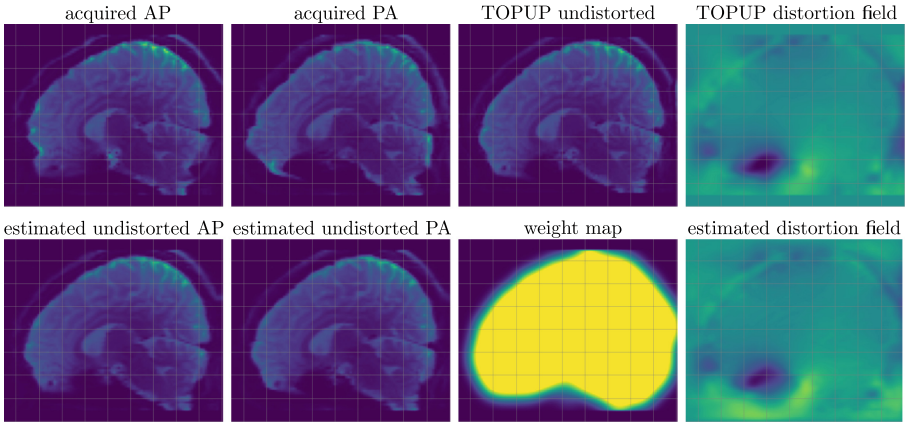


**Fig. 3.** Fidelity (MSE) to ground truth distortion fields and images for the different models for real data with TOPUP as ground truth and for synthetic data. The field fidelity is expressed as a fraction of the voxel dimension (2 mm). Large frames share the same scale for global overview between experiments whereas yellowish ones are zoomed in for intra-metric, between models comparison. (Color figure online)

- The real dataset experiment uses real acquired EPI images and evaluate how close to TOPUP (ground truth) the different models behave. In terms of field fidelity, supervised and semi-supervised models show a mean voxelique MSE around 0.2 (equivalent to  $0.4\text{mm}^2$ ), way better than the unsupervised approach that shows an error about 4 times bigger with also a much higher variance. In terms of image similarity, the unsupervised and the semi-supervised approach show an error equivalent to 2/3 of the one of the supervised approach with akin variance. The semi-supervised approach performs well for both metrics whereas the other models present weaker results in one situation. An example case subject from the testing sample, corrected with the semi-supervised approach, is shown in Fig. 4.
- The synthetic dataset experiment uses simulated EPI images and evaluate how each model and TOPUP are able to retrieve a synthetic ground truth. It quite follow the same trend as above for the deep-learning models. TOPUP, the supervised and the supervised models show similar field fidelity that is better than the one of the unsupervised model. TOPUP, the unsupervised and the supervised models show similar image fidelity that is better than the one of the supervised model. TOPUP and the semi-supervised model performs similarly well for both metrics whereas the supervised and the unsupervised models are weaker in one situation.

**Table 1.** Fidelity (MSE: mean (std)) to ground truth distortion fields and images for the different models for real data with TOPUP as ground truth and for synthetic data. The field fidelity is expressed as a fraction of the voxel dimension (2 mm).

		Uncorrected	TOPUP	Unsupervised	Supervised	Semi-supervised
Real dataset	Field fidelity	2.21 (0.70)	$\emptyset$	0.86 (0.34)	0.21 (0.07)	0.22 (0.07)
	Image fidelity	$1.39 \cdot 10^3$ ( $3.39 \cdot 10^4$ )	$\emptyset$	$2.22 \cdot 10^4$ ( $6.76 \cdot 10^5$ )	$3.27 \cdot 10^4$ ( $9.09 \cdot 10^5$ )	$2.10 \cdot 10^4$ ( $6.29 \cdot 10^5$ )
Synthetic dataset	Field fidelity	2.21 (0.70)	0.17 (0.02)	0.36 (0.11)	0.17 (0.05)	0.21 (0.08)
	Image fidelity	$3.89 \cdot 10^3$ ( $5.94 \cdot 10^4$ )	$8.73 \cdot 10^4$ ( $1.51 \cdot 10^4$ )	$8.98 \cdot 10^4$ ( $1.82 \cdot 10^4$ )	$1.55 \cdot 10^3$ ( $3.41 \cdot 10^4$ )	$9.25 \cdot 10^4$ ( $1.81 \cdot 10^4$ )



**Fig. 4.** Example case from the testing sample (unseen) of the real dataset, corrected using the semi-supervised model.

## 6 Discussion

This study shows that, with proper distortion model, one can use deep-learning registration trained on some tens of subjects to rapidly correct a larger set for susceptibility-induced distortion, with results as good as TOPUP. One can use a processed subset of a dataset to include transformation supervision and improve the transformation fidelity which is a direct measure of the quality of the registration. One might have expected the semi-supervised model to perform half-way between its unsupervised and supervised counterparts in terms of field and image fidelity. However, our findings suggest it offers the best of both worlds.

Similar to TOPUP [4], the proposed approach requires reversed PED acquisitions to correct distortion in EPI images through registration, although this process is performed using a deep-learning architecture, much faster than its classical counterpart. Our approach builds on the purely unsupervised approach

in [8], but enhance the training notably using more reliable transformation supervision.

In the broader context of applying deep learning to EPI distortion correction, Schilling et al. [3] proposes a technique that allows TOPUP-based correction in the absence of a second reversed PED acquisition. It works by applying deep learning-based modality transfer to synthesise a distortion-free  $b = 0$  from a T1-weighted scan; the distortion-free  $b = 0$  is then used with the acquired  $b = 0$  for TOPUP processing. Our method can be readily extended to drastically shorten this distortion correction process as well. It could also be used as the primary correction prior to the FOD-based secondary one from [5].

In future work, we plan to evaluate the effect of the training sample size on performance. We also intend to study the generalizability of the proposed approach. Although decent performance may be expected when training and testing on different datasets, due to the relative homogeneity of  $b = 0$  images compared to anatomical ones for example, the high variability of acquisition parameters (TE, TR...) may still lead to sub-optimal results.

## 7 Conclusion

We presented a semi-supervised approach for the distortion correction of EPI images with opposite PED using DL-based image registration. We compared this model with an unsupervised and a supervised one, as well as a traditional algorithm: TOPUP. By leveraging ground truth distortion transformations during training, the proposed method can produce more accurate estimate of distortion fields (direct quantitative metric) compared to unsupervised approaches at testing. It also outperforms the supervised approach for the image metric. On synthetic data, the results were similar to TOPUP but with much faster computation. The proposed model can typically be trained on a processed subsample of a dataset with an external tool and then be applied to the rest of the dataset to produce distortion correction very efficiently.

## References

1. Graham, M.S., Drobnjak, I., Jenkinson, M., Zhang, H.: Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI. *PLoS One* **12**(10), e0185647 (2017)
2. Gu, X., Eklund, A.: Evaluation of six phase encoding based susceptibility distortion correction methods for diffusion MRI. *Front. Neuroinform.* **13**, 76 (2019)
3. Schilling, K.G., et al.: Distortion correction of diffusion weighted MRI without reverse phase-encoding scans or field-maps. *PLoS One* **15**(7), e0236418 (2020)
4. Andersson, J.L., Skare, S., Ashburner, J.: How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage* **20**(2), 870–888 (2003)
5. Qiao, Y., Shi, Y.: Unsupervised deep learning for FOD-based susceptibility distortion correction in diffusion MRI. *IEEE Trans. Med. Imaging* **41**(5), 1165–1175 (2022)

6. Begnoche, J.P., Schilling, K.G., Boyd, B.D., Cai, L.Y., Taylor, W.D., Landman, B.A.: EPI susceptibility correction introduces significant differences far from local areas of high distortion. *Magn. Reson. Imaging* **92**, 1–9 (2022)
7. Zahneisen, B., Baumler, K., Zaharchuk, G., Fleischmann, D., Zeineh, M.: Deep flow-net for EPI distortion estimation. *Neuroimage* **217**, 116886 (2020)
8. Duong, S.T.M., Phung, S.L., Bouzardoum, A., Schira, M.M.: An unsupervised deep learning technique for susceptibility artifact correction in reversed phase-encoding EPI images. *Magn. Reson. Imaging* **71**, 1–10 (2020)
9. Rohlfing, T.: Image similarity and tissue overlaps as surrogates for image registration accuracy: widely used but unreliable. *IEEE Trans. Med. Imaging* **31**(2), 153–163 (2012)
10. Drobnjak, I., Gavaghan, D., Süli, E., Pitt-Francis, J., Jenkinson, M.: Development of a functional magnetic resonance imaging simulator for modeling realistic rigid-body motion artifacts. *Magn. Reson. Med.* **56**(2), 364–380 (2006)
11. Jezzard, P., Balaban, R.S.: Correction for geometric distortion in echo planar images from B0 field variations. *Magn. Reson. Med.* **34**(1), 65–73 (1995)
12. Morgan, P.S., Bowtell, R.W., McIntyre, D.J., Worthington, B.S.: Correction of spatial distortion in EPI due to inhomogeneous static magnetic fields using the reversed gradient method. *J. Magn. Reson. Imaging* **19**(4), 499–507 (2004)
13. Chang, H., Fitzpatrick, J.M.: A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE Trans. Med. Imaging* **11**(3), 319–29 (1992)
14. Fu Y., Lei Y., Wang T., Curran W.J., Liu T., Yang X.: Deep learning in medical image registration: a review. *Phys. Med. Biol.* **65**, 20TR01 (2020)
15. Haskins, G., Kruger, U., Yan, P.: Deep learning in medical image registration: a survey. *Mach. Vis. Appl.* **31**, 8 (2020)
16. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J.V., Dalca, A.V.: Voxel-Morph: a learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* **38**, 1788–1800 (2019)
17. Dalca, A.V., Balakrishnan, G., Guttag, J.V., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Med. Image Anal.* **57**, 226–236 (2019)
18. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)