EPI Distortion Correction without Opposite Phase Encodings with Unsupervised INR-Based Deformable Registration

Tyler Spears[1] iD   TAS6HH@VIRGINIA.EDU
[1] *Dept. of Electrical & Computer Engineering, University of Virginia, Charlottesville, VA, USA*

Myla Goldman[2] iD   MYLA.GOLDMAN@VCUHEALTH.ORG
[2] *Dept. of Neurology, Virginia Commonwealth University School of Medicine, Richmond, VA, USA*

P. Thomas Fletcher[1] iD   PTF8V@VIRGINIA.EDU

Abstract

Diffusion MRIs (dMRIs) provide a detailed look at the structure of the brain, but the acquired images come with many distortions. Echo planar imaging (EPI) sequences, nearly universal for dMRIs, are highly sensitive to inhomogeneities of the magnetic field in the scanner. This results in severe geometric distortion (up to tens of millimeters) in the phase encoding direction, particularly in areas with strong changes in tissue density such as the brainstem, temporal, and frontal regions. A common method for correcting EPI distortion is to collect an image with the opposite phase encoding (PE) direction and reconstruct the magnetic susceptibility field. However, many dMRI protocols, some still in use today, do not include this auxiliary acquisition. Other methods have attempted to register the distorted EPI to an anatomical reference, with less accurate results. In this work, we propose EPINR, an unsupervised implicit neural representation (INR) based registration model that builds on these previous works. EPINR learns the susceptibility field by warping a single b0 image to a T1 reference, without opposite PE acquisitions. EPINR also leverages its smooth and continuous representation to apply higher-order regularizations calculated analytically. We evaluate EPINR against several comparison methods, both traditional and learning-based, over two dMRI datasets. We then discuss the reasons for EPINR's high performance, and how it can bring structural precision to previously compromised diffusion images.

Keywords: Intermodal Registration, INR, Susceptibility Distortion Correction

1. Introduction

Diffusion MRI (dMRI) is the most detailed structural description of the brain we can acquire *in vivo*. Unfortunately, that detail is compromised by severe geometric distortions, on the order of tens of millimeters, caused by magnetic field inhomogeneities in the scanner (Farzaneh et al., 1990; Haskell et al., 2023). These distortions affect nearly all dMRI scan sequences, and field inhomogeneities are produced whenever any tissue or material is in the scanner. So, nearly all dMRIs need this distortion, often called susceptibility distortion or echo planar image (EPI) distortion, to be corrected.

Several susceptibility distortion correction (SDC) methods have been previously proposed, but the most effective methods require auxiliary acquisitions explicitly for SDC. For

example, the popular FSL `topup` (Andersson et al., 2003) tool performs SDC with impressive accuracy (Gu and Eklund, 2019; Graham et al., 2017; Wang et al., 2017), and is a staple in dMRI processing pipelines (Glasser et al., 2013; Cieslak et al., 2021). The `topup` tool requires only a single b0 collected with the opposite phase encoding (PE) direction as the primary diffusion weighted images (DWIs). However, that single image is not collected in many dMRI datasets, especially in places that rely on older sequences. Examples include: the Open Access Series of Imaging Studies 3 (OASIS3) dataset (LaMontagne et al., 2019), the Center on Reproducible Research (CoRR) dataset (Zuo et al., 2014), the MASiVar multisite and multiscanner dataset (Cai et al., 2021), the TractoInferno dataset (Poulin et al., 2021), and many other public (and private) dMRI datasets.

Alternative SDC methods have been proposed that only require an undistorted anatomical reference image, such as a T1-weighted (T1w) or T2-weighted (T2w) image, which are almost always collected in dMRI protocols (Kybic et al., 2000). However, these deformable registration-based methods are significantly less accurate than methods, such as `topup`, that require auxiliary images (Gu and Eklund, 2019; Graham et al., 2017; Wang et al., 2017). New unsupervised implicit neural representation (INR) (Sitzmann et al., 2020) registration methods (Wolterink et al., 2022; Byra et al., 2023) could bring new life into this decades-old problem of SDC without extra acquisitions.

In this work, we propose EPINR, an unsupervised INR-based deformable registration model for EPI distortion correction without extra acquisitions. EPINR learns a smooth displacement field, constrained to the PE direction, that aligns a b0 and T1 image pair for a particular subject. EPINR also leverages the spatially-continuous and smooth nature of INRs to *analytically* compute Jacobians and higher-order derivatives for regularization (Rueckert et al., 1999; Wolterink et al., 2022). Our contributions in this work are as follows:

- EPINR, an unsupervised INR-based registration model that is the first of its kind to be applied to SDC without auxiliary acquisitions.

- A comparison of previous methods, traditional and deep learning-based, on two DWI datasets with vastly different image characteristics, including an in-house dataset not used in any previous work on SDC validation.

- Validation that EPINR outperforms previous methods, and an analysis of where a previous state of the art method struggles.

2. Background

EPI distortions cause a strong geometric warping during the data readout step of an EPI sequence, which effects both dMRIs and functional MRIs (fMRIs) (Studholme et al., 2000; Gholipour et al., 2006). This is caused by inhomogeneities in the magnetic susceptibility field in the scanner, causing a non-linear shift in the scanner's gradient field at tissue boundaries (Johansen-Berg and Behrens, 2014, Chapter 4). This distortion is almost entirely constrained to the phase encoding (PE) direction, and it is most severe in the sinuses and the brainstem, while also affecting temporal and frontal regions (Treiber et al., 2016). We also note that for dMRIs, SDC is usually performed only on images with no diffusion weighting (b0s), as they do not contain eddy current-induced distortions (Andersson et al., 2017; Rohde et al., 2004).

2.1. SDC with Extra Acquisitions

As mentioned, SDC methods that leverage auxiliary acquisitions often reconstruct the most accurate susceptibility fields (Gu and Eklund, 2019; Graham et al., 2017; Wang et al., 2017). This includes fieldmap-based SDC and reverse PE SDC. Some methods suggested collecting fieldmaps that measure the magnetic inhomogeneity of the subject (Jezzard and Balaban, 1995; Wan et al., 1997; Reber et al., 1998), similar to those in fMRI, but comparisons (Tao et al., 2009; Fritz et al., 2014; Wang et al., 2017) have found that these approaches are sensitive to acquisition parameters and partial volume effects. Reverse PE SDC is more common, where a b0 image with an opposite PE direction is collected. This only adds a few seconds to the scan time, but is still absent in many datasets and sequences. That said, reverse PE images are sufficiently common, and their performance is very high, such that they are raised to a "silver standard" in SDC (Gu and Eklund, 2019). These methods assume that images with opposite PEs will have an equal and opposite displacement field, and a parameterized field is optimized with these constraints.

Reverse PE SDC was first proposed by Chang and Fitzpatrick (1992) and popularized with the FSL `topup` tool (Andersson et al., 2003). Other methods of this class continue to be developed, such as those in Holland et al. (2010), Irfanoglu et al. (2015) (DR-BUDDI), Hédouin et al. (2017), and Liu et al. (2021). Deep learning has also been applied here, with EPIs of both PE directions given to a convolutional neural network (CNN) for predicting the susceptibility field (Hu et al., 2020; Qiu et al., 2026; Zahneisen et al., 2020; Legouhy et al., 2022; Qiao and Shi, 2022; Zaid Alkilani et al., 2024). We also mention that pulse sequence design for reducing EPI distortion is an active area of research (Haskell et al., 2023), but these sequences are still experimental and cannot help with dMRIs that are already acquired.

2.2. Deep Learning Deformable Registration

VoxelMorph & CNN Registration. Deep learning-based registration methods have exploded in popularity after the popular VoxelMorph framework (Balakrishnan et al., 2019) was proposed. VoxelMorph uses a CNN that takes the fixed and moving images as inputs and predicts a deformation field that warps the moving image to the fixed. The primary objective function is an image similarity measure, but also includes regularization terms that penalize the squared norm of the displacement spatial gradient. Note that this gradient is approximated with a finite difference method.

INR-Based Registration. INRs, first proposed as sinusoidal representation networks (SIRENs) in Sitzmann et al. (2020), are multilayer perceptrons (MLPs) that learn a continuous signal representation by mapping coordinates to intensity values. INR registration represents a warp field by using an image similarity objective function, rather than image reconstruction (Wolterink et al., 2022). If properly constructed, this representation is C^∞ smooth, and may analytically calculate higher order derivatives and use advanced regularization techniques (Rueckert et al., 1999). Wolterink et al. (2022) was one of the first of these methods with implicit deformable image registration (IDIR). Sun et al. (2024) proposed the Neural Image Registration (NIR) model, which extended IDIR with a hybrid coordinate sampling scheme, and van Harten et al. (2024) proposed a cycle-consistent de-

formation with dual INRs. Byra et al. (2023) analyzed the impact of different INRs on deformable brain image registration.

2.3. Previous Works in SDC Without Auxiliary Acquisitions

Traditional Registration. A third class of SDC method is to register the distorted b0 to an undistorted anatomical reference, constrained to the PE direction. While this problem may seem straightforward, the intermodal differences between a b0 (T2-weighted) and a T1w, along with the difference in noise levels and resolution, prove challenging. Kybic et al. (2000) was the first to propose a deformable registration SDC method. Later, Gholipour et al. (2006) proposed a Free-Form Deformation method for fMRIs, and Tao et al. (2009) used a variational approach to register the b0 with a *T2w* reference. However, one of the most common models for this approach to apply a general image registration framework, such as ANTs symmetric normalization (SyN) (Avants et al., 2008) to the task, as shown in Wang et al. (2017); Gu and Eklund (2019). Unfortunately, registration-based correction, while requiring no extra acquisitions, also gives the lowest performance of the three SDC methods (Gu and Eklund, 2019; Graham et al., 2017; Wang et al., 2017).

Learning-Based Registration The availability of large DWI repositories has inspired applications of neural networks to EPI distortion correction. As mentioned previously, many works have used the opposite PE approach to predict the susceptibility fields, but few have attempted SDC without auxiliary images. Ye et al. (2021) proposed a model to perform simultaneous super-resolution and distortion correction with only one PE direction. Jimeno et al. (2024) proposed GDCNet, which applies VoxelMorph to EPI distortion in fMRIs, but has not yet been peer-reviewed at the time of this writing. Possibly the most popular deep network SDC method is Synb0 (Schilling et al., 2019), which is a conditional generative adversarial network (GAN) that synthesizes an undistorted b0 when given a distorted b0 and a T1. This method was trained in 2.5D (2D slices in all three axes) on a dataset of over 500 subjects, with output images at 2.5mm isotropic. The real and generated b0 images are then given to `topup` for SDC. This model was used for SDC for datasets in Begnoche et al. (2022), Poulin et al. (2021),Cai et al. (2021), and more.

3. EPINR

Our goal with EPINR is to correct for EPI distortion in a diffusion image protocol using only a single b0 (effected by EPI distortion) and a T1-weighted anatomical reference (not effected by this distortion), and our model is illustrated in Figure 1. The underlying susceptibility field can be calculated simply by scaling a displacement field $u(x, y, z)$, so registration and field reconstruction are the same problem here. Like classical registration models, we train EPINR in an unsupervised manner with a newly initialized model for every subject.

Suppose we have a spin echo pulse sequence (Johansen-Berg and Behrens, 2014, Chapter 2) that produces a b0 image $\mathbf{M}$ distorted only in the PE direction, and a non-distorted T1 image $\mathbf{F}$. We assume, without loss of generality, that the PE direction is y. Additionally, suppose that as a preprocessing step, $\mathbf{F}$ has been rigidly aligned to match the orientation and translation of $\mathbf{M}$. Note that we keep $\mathbf{F}$ and $\mathbf{M}$ in their native resolutions as we will just sample locations in $\mathbf{F}$ and $\mathbf{M}$ during training. We denote this sampling operation as $\mathbf{A}[\mathbf{c}]$, which is a trilinear interpolation of volume $\mathbf{A}$ at the 3D physical coordinate $\mathbf{c}$. So,
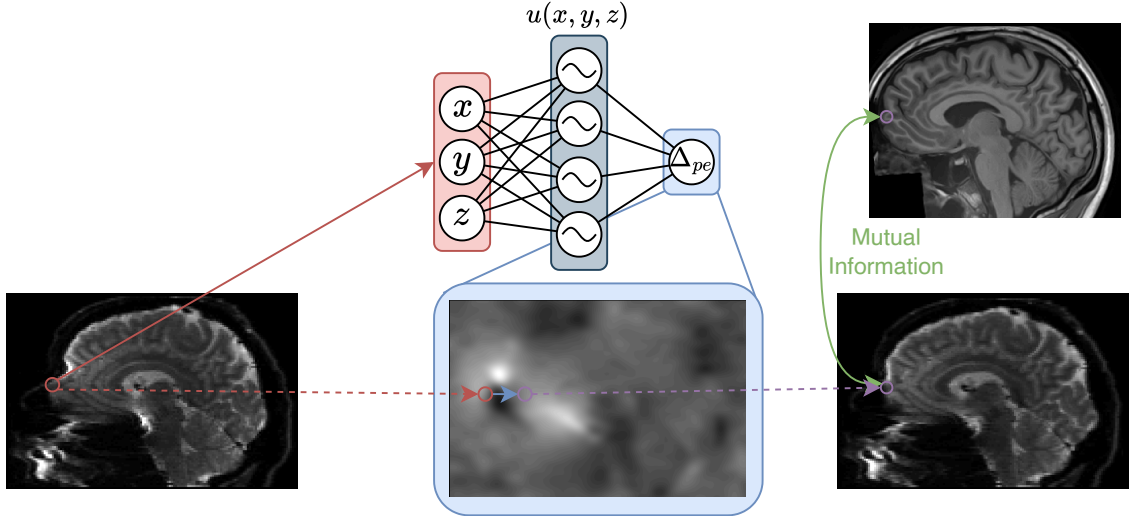
Figure 1: Illustration of our proposed EPINR registration model. EPINR learns to warp a distorted b0 to a non-distorted T1 reference in an unsupervised manner with a displacement field (blue) constrained in the phase encoding direction. The C^∞ smoothness of EPINR also allows for *analytically* calculating higher-order regularization methods and Jacobian modulation (Jezzard and Balaban, 1995).

our network, a SIREN (Sitzmann et al., 2020) in our implementation, $u_\theta : \mathbb{R}^3 \to \mathbb{R}$ with parameters θ, maps $u_\theta(x_i, y_i, z_i) = (0, \Delta y_i, 0)$, with zeros in the non-PE directions.

We now describe two domain-specific features of EPINR. First, it is known that EPI distortion in spin echo sequences is mass preserving (Chang and Fitzpatrick, 1992; Jezzard and Balaban, 1995) in that compressed signal is "piled-up" to a higher intensity, and vice-versa for "stretched" areas. This is expressed as a division of the image intensity by the determinant of the Jacobian of the deformation field, and we need to compensate for that when applying our predicted displacement field. This is also known as Jacobian modulation (Andersson et al., 2003). Second, while the b0 and T1 reference are rigidly aligned during preprocessing, the transformation cannot be fully determined before SDC. So, we define an additional set of trainable parameters $\rho \in \mathbb{R}^6$ that form a rigid transformation A_ρ. This is applied to the T1 during training so that any rigid components of u_θ may be discouraged.

Our similarity loss is defined as:

$$\mathcal{L}_{\text{sim}} = \frac{1}{N} \sum_{i=1}^{N} -\text{MI}\left(\det J_u\left(\mathbf{c}_i + u_\theta(\mathbf{c}_i)\right) \mathbf{M}\left[\mathbf{c}_i + u_\theta(\mathbf{c}_i)\right], \mathbf{F}\left[A_\rho \mathbf{c}_i\right]\right), \tag{1}$$

where $\mathbf{c}_i$ is a patch of contiguous coordinates in the i'th batch, J_u is the Jacobian of the deformation field produced by the EPINR network u, and MI is the mutual information similarity (Viola and Wells III, 1997). The $\det J_u$ term applies Jacobian modulation, but this modulation is changing along with the network weights during training. This complicates the optimization process (as mentioned in Andersson et al. (2001); Tao et al. (2009)), but

potentially increases final registration accuracy as modulation will always be applied to the final undistorted b0. We also reiterate that J_u is analytically calculated on the network parameters θ, not with a finite difference approximation (Wolterink et al., 2022).

Taking inspiration from the SDC method in fmriprep (Esteban et al., 2019), we also added a similarity term for the Laplacians of the $\mathbf{F}$ and the undistorted $\mathbf{M}$:

$$\mathcal{L}_{\text{lap}} = \frac{1}{N} \sum_{i=1}^{N} -\text{NCC}\left(\nabla^2\left(\det J_u\left(\mathbf{c}_i + u_\theta(\mathbf{c}_i)\right)\mathbf{M}\left[\mathbf{c}_i + u_\theta(\mathbf{c}_i)\right]\right), \nabla^2(1 - \mathbf{F}\left[A_\rho \mathbf{c}_i\right])\right), \quad (2)$$

where NCC is normalized cross-correlation, ∇^2 is the image Laplacian operator, and $1 - \mathbf{F}$ flips the T1 intensity distribution to more closely match the b0.

Finally, we used two regularization functions: a smoothing loss as the Frobenius norm of the Jacobian, and a bending energy loss that requires the Hessian of the deformation (Rueckert et al., 1999). These are given as:

$$\mathcal{L}_{\text{smooth}} = \|J_u(\mathbf{c}_i + u_\theta(\mathbf{c}_i))\|_F = \left|1 + \frac{\partial u_\theta(\mathbf{c}_i)}{\partial y}\right| \quad (3)$$

$$\mathcal{L}_{\text{bend}} = \sum_{p=1}^{3}\sum_{q=1}^{3}\sum_{r=1}^{3}\left(\mathbf{H}_p^{(u)}(\mathbf{c}_i + u_\theta(\mathbf{c}_i))\right)_{q,r}^2 = \sum_{q=1}^{3}\sum_{r=1}^{3}\left(\mathbf{H}_{p=2}^{(u)}(\mathbf{c}_i + u_\theta(\mathbf{c}_i))\right)_{q,r}^2, \quad (4)$$

with simplified forms on the right side enabled by the constraint to the PE direction. Here, $(\mathbf{H}_p^{(u)}(\mathbf{c}_i + u_\theta(\mathbf{c}_i))_{q,r}$ is the analytically calculated Hessian matrix of the deformation field produced by the network u_θ at the p'th output dimension with relation to input dimensions q and r (where $x = 1$, $y = 2$, and $z = 3$).

Together with their corresponding weight terms, these four loss terms form our full objective function. After training, we discard the learned ρ as it is only meant to account for any rigid misalignment that we want the displacement field to ignore during training. In practice, A_ρ was frozen halfway during training, and only ever contributed a minuscule amount of movement.

4. Experiments & Results

4.1. Data

We evaluated EPINR and the comparison methods on two different datasets. All methods were either unsupervised or trained on other datasets, so there were no train/test splits. All evaluation b0s were acquired with an anterior to posterior (AP) PE direction, and only one b0 was used. All DWIs were denoised with Marchenko-Pastur principal component analysis (MP-PCA) (Veraart et al., 2016) and had B_1 bias field correction applied with N4 (Tustison et al., 2010). For topup, one AP b0 and one PA b0 was used for SDC. T1 images were bias corrected with N4 and denoised with a non-local means filter (Manjón et al., 2010). Then, T1 images were rigidly registered to the *distorted* b0s with ANTs neighborhood normalized CC (Avants et al., 2008). Tissue masks were estimated with the FSL brain extraction tool (BET) (b0) (Smith, 2002) and SynthStrip (T1) (Hoopes et al., 2022). These masks were used to weigh voxels during training.

6

The first dataset was the publicly available Microstructure-Informed Connectomics (MICA-MICS) dataset (Royer et al., 2021) with $N = 49$ subjects. DWIs were collected with a spin-echo sequence at 1.6mm isotropic resolution, TR=3500ms, TE=64.40ms, and an FSL-style total readout time of 0.05282 seconds. T1 images were acquired with an MP-RAGE sequence, 0.8mm isotropic resolution, TR=2300ms, and TE=3.14ms. The second dataset was an in-house DWI dataset collected at Virginia Commonwealth University (VCU) as part of a study involving multiple sclerosis (MS) (Pearsall et al., 2024; Goldman et al., 2023), which we name "VCU MS." This study was approved by the VCU Institutional Review Board. We evaluated methods on $N = 48$ patients from VCU MS, including 10 healthy controls and 38 patients diagnosed with MS. DWIs were acquired at 1.875mm in-plane and 2mm slice thickness, TR=4900ms, TE=118ms, and a total readout time of 0.063054 seconds. The T1 images were acquired at 1mm isotropic resolution with TE=2.99ms.

4.2. Comparison Methods

We compare EPINR to several SDC methods on both datasets:

1. *EPINR* - Our proposed model that uses an INR to learn the displacement field between a b0 and a T1 image, without resizing either out of their native resolutions.

2. *EPINR-128* - An ablated EPINR with only 4 layers and 128 hidden features.

3. *Uncorrected* - The acquired b0 with preprocessing and no SDC applied.

4. *ANTs-SyN* - ANTs Syn deformable registration (Avants et al., 2008) with parameters from the `antsIntermodalityIntrasubject.sh` script with PE direction constraints.

5. *QSIPrep* - Fieldmapless SDC built into the QSIPrep dMRI preprocessing pipeline (Cieslak et al., 2021), version 1.0.2. This is a variant of ANTs-SyN (Avants et al., 2008) with tweaked preprocessing and registration parameters; this is designated as "experimental" within QSIPrep.

6. *Synb0+Topup* - Synb0 synthesized undistorted b0 volume with follow-up SDC with topup (Schilling et al., 2019). Synb0 uses a trained GAN to synthesize an undistorted b0 at 2.5mm isotropic resolution, given a masked distorted b0 and a masked T1. The preprocessed b0s and T1s were masked for input, using the scilpy (Boré et al., 2025) wrapper around a Synb0 implementation. The synthetic b0 is then upsampled to the b0's native resolution, and topup performs SDC.

7. *Topup GT* - FSL topup (Andersson et al., 2003) run with b0s that have AP and PA PE directions, considered as our ground truth.

4.3. Implementation Details

We implemented EPINR as a SIREN (Sitzmann et al., 2020) MLP with 256 hidden features and 5 layers. Training was performed on randomly sampled patches as a full b0 could not fit into GPU memory, effective batch size of 6, 50 batches per epoch, over 60 epochs. Adam with weight decay (Loshchilov and Hutter, 2018) was used for optimization, with an initial

learning rate of 5×10^{-4} that linearly decreased to 5×10^{-6} with a small number of warmup and cooldown epochs. Both the b0 and T1 volumes were normalized to $[0, 1]$. The input coordinates $\mathbf{c}_i$ were normalized to $[-1, 1]$ based on the b0's field of view, and MLP outputs were re-scaled to physical units. The image Laplacian ∇^2 in Equation 2 was implemented as a difference of Gaussian filters (Marr and Hildreth, 1980).

4.4. Evaluation Metrics

We transform each model's predicted displacement field (in mm) to a susceptibility field (Hz) with scaling by the PE direction resolution and total readout time. Then, we use the `applytopup` FSL tool to apply the susceptibility field to the distorted b0 with Jacobian modulation. We then measure each model's performance over five metrics: 1) a mean-squared error (MSE) of the model's undistorted b0 vs. topup's undistorted b0, 2) MSE of the model's predicted susceptibility field vs. topup's reconstructed field, 3) MI similarity between the model's corrected b0 and the T1 reference, with both normalized to $[0, 1]$, 4) local normalized CC (LNCC) of the corrected b0 and the T1 reference image, and 5) percentage of negative voxels in the determinant of the Jacobian of the model's deformation field, calculated with a central difference method to keep evaluation consistent.

4.5. Results

Dataset	Model	b0-topup MSE $\times 10^{-6}$ ↓	SF-topup MSE $\times 10^{-3}$ ↓	MI b0-Anat. ↑	LNCC b0-Anat. ↑
MICA-MICS	Uncorrected	1.3021 (0.300)	0.5397 (0.125)	0.3675 (0.034)	0.2271 (0.014)
	ANTs-SyN	1.6665 (0.354)	0.5229 (0.092)	0.3342 (0.054)	0.2395 (0.017)
	QSIPrep	1.4983 (0.361)	0.5710 (0.128)	0.3677 (0.033)	0.2252 (0.012)
	Synb0+Topup	1.3741 (0.338)	1.4491 (0.267)	0.4285 (0.034)	0.2138 (0.013)
	EPINR-128	**0.9247 (0.221)**	0.4827 (0.102)	0.4107 (0.045)	0.2429 (0.016)
	EPINR	0.9553 (0.226)	**0.4247 (0.084)**	**0.4306 (0.048)**	**0.2482 (0.016)**
	Topup GT	-	-	0.4629 (0.039)	0.2639 (0.019)
VCU MS	Uncorrected	4.1649 (1.375)	0.8881 (0.248)	0.2750 (0.040)	0.1621 (0.016)
	ANTs-SyN	3.5424 (1.448)	0.7522 (0.189)	0.2167 (0.039)	0.1840 (0.016)
	QSIPrep	3.2523 (1.066)	0.8075 (0.222)	0.2667 (0.038)	0.1720 (0.017)
	Synb0+Topup	3.3660 (1.324)	2.7373 (0.900)	0.3276 (0.023)	0.1540 (0.016)
	EPINR-128	2.2986 (1.853)	0.6500 (0.557)	0.3334 (0.029)	0.1832 (0.018)
	EPINR	**2.1254 (1.801)**	**0.6234 (0.531)**	**0.3666 (0.030)**	**0.1866 (0.018)**
	Topup GT	-	-	0.3146 (0.027)	0.1780 (0.020)

Table 1: Results of different EPI distortion correction methods applied to two DWI datasets. Metric name arrows indicate the direction of better performance, and the best performance between all models is bolded.

Our results for all datasets, models, and metrics (except the Jacobian negative percentage) are shown in Table 1, and examples from both datasets are shown in Figure 2. For the negative Jacobian percentage, the large majority of methods showed 0 negative voxels, except for: Synb0+topup on MICA-MICS (0.0152%), Synb0+topup on VCU MS (0.0017%), and regular topup on VCU MS (0.0006%).
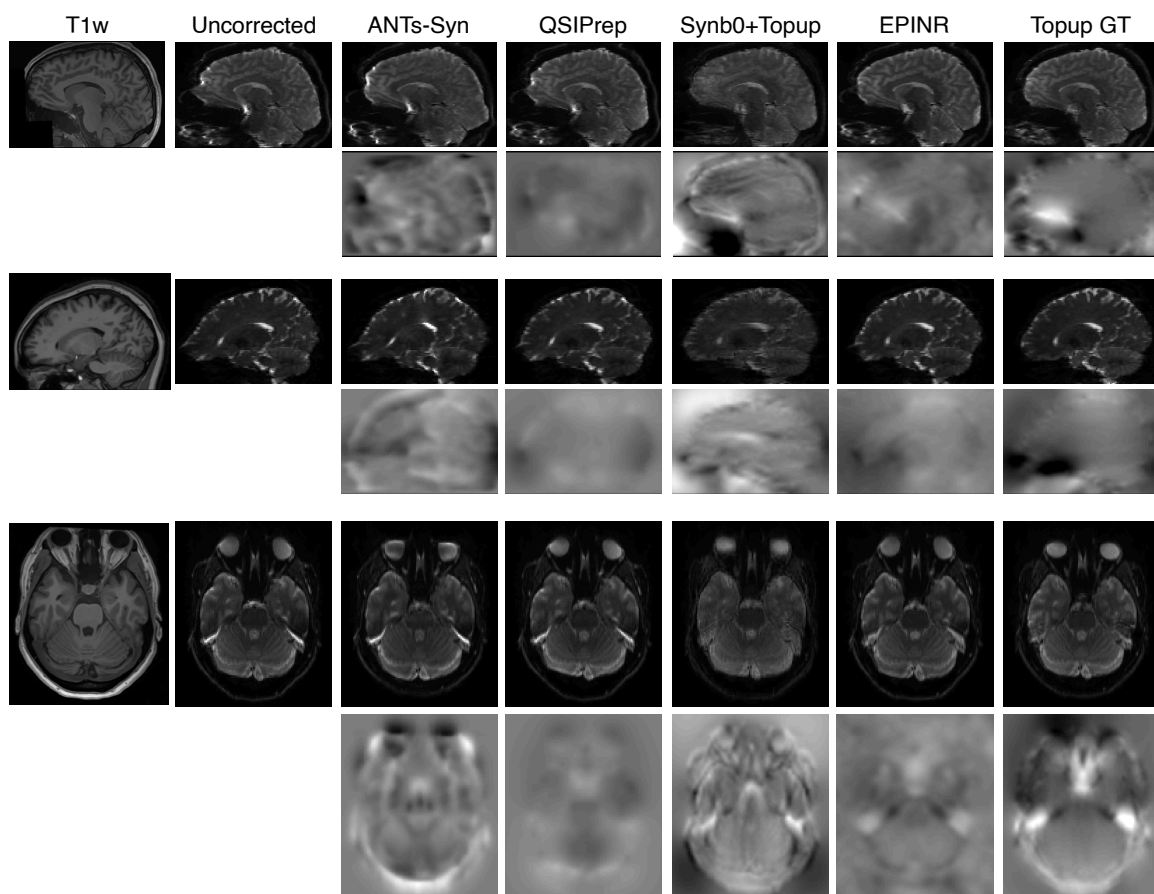
Figure 2: Qualitative results of susceptibility distortion correction over three subjects. Below each undistorted b0, the predicted susceptibility field in Hz is displayed. From top to bottom, subjects are from MICA-MICS, VCU MS, and MICA-MICS.

As shown in Table 1, EPINR outperforms comparison methods on every metric we evaluated over both datasets except for b0 MSE in the MICA-MCIS data, where the ablation model performed the best. We believe that this overall high performance may be due to the high number of parameters that allows EPINR to learn high-frequency warps. According to Figure 2 rows 1 and 3, we also see how including the Jacobian modulation in training helps EPINR "break through" some of the more severe warping in frontal areas, compared to ANTs-SyN and QSIPrep. The 0 negative Jacobian voxels also indicate that the smooth and bending energy regularizers are effective at eliminating "folds" in the resulting deformation fields, in agreement with previous work in Wolterink et al. (2022).

The performance of Synb0 is somewhat surprising. Despite its use in several publicly released datasets, Synb0 performed the worst on some of our metrics. The predicted susceptibility fields in Figure 2 indicate that the topup registration to the synthesized b0 is a poor fit. This may be explained by the fixed spatial resolution of 2.5mm isotropic generated by Synb0, which must be upsampled to 1.6mm for MICA-MICS and 1.875mm for VCU MS and may be too large of a resolution gap for topup to handle. We also found that the generated b0s "filled out" brain tissue into the outer surface of the skull, which changed the basic shape and volume of the b0. This may be caused by an unexpected brain mask shape resulting from our preprocessing, but this would imply that Synb0 is overly sensitive to how masking is performed, which is a difficult task in distorted images.

The results also reveal the limitations of calculating image similarity between the b0 and T1 reference. In the VCU MS data, EPINR "outperforms" topup on image similarity for both MI and LNCC. However, a qualitative comparison of topup and EPINR still confirms that topup produces a better undistorted b0. This discrepency is likely due to 1) EPINR's much higher number of parameters allowing it to somewhat overfit to the anatomical reference, and 2) the simple fact that image similarity between two modalities is not on its own sufficient to measure SDC performance. So, while b0-T1 similarity is often reported in the literature, interpreting results must be done carefully, especially in the presence of over-parameterized deep learning models.

5. Discussion

We have proposed EPINR, a novel INR-based unsupervised registration method for correcting EPI distortions in dMRIs without auxiliary acquisitions. We compared EPINR to several models currently in use, and found that EPINR produced higher quality susceptibility field reconstructions. We also validated EPINR's flexibility over two dMRI datasets with different resolutions and image characteristics.

We believe that EPINR's flexibility lends itself to further extensions. While EPINR's unsupervised nature is appealing, we believe that performance could be greatly improved by incorporating learned priors into the model via methods such as meta-learning (Finn et al., 2017). We would also be interested in expanding our analyses to look at the downstream effects of these SDC methods on microstructure modeling and tractography. Finally, we would also like to expand EPINR to EPI corrections in fMRI, although this task is more challenging as distortion in gradient echo images causes signal dropout instead of "pile-up." If this could be tackled, EPINR would continue to let us to make older data more useful.

Acknowledgments

References

Jesper L. R. Andersson, Stefan Skare, and John Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *NeuroImage*, 20(2):870–888, October 2003. ISSN 1053-8119.

Jesper L. R. Andersson, Mark S. Graham, Ivana Drobnjak, Hui Zhang, Nicola Filippini, and Matteo Bastiani. Towards a comprehensive framework for movement and distortion correction of diffusion MR images: Within volume movement. *NeuroImage*, 152:450–466, May 2017. ISSN 1053-8119.

Jesper L.R. Andersson, Chloe Hutton, John Ashburner, Robert Turner, and Karl Friston. Modeling Geometric Deformations in EPI Time Series. *NeuroImage*, 13(5):903–919, May 2001. ISSN 10538119.

B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1):26–41, February 2008. ISSN 1361-8415.

Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, and Adrian V. Dalca. VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *IEEE Transactions on Medical Imaging*, 38(8):1788–1800, August 2019. ISSN 1558-254X.

John P. Begnoche, Kurt G. Schilling, Brian D. Boyd, Leon Y. Cai, Warren D. Taylor, and Bennett A. Landman. EPI susceptibility correction introduces significant differences far from local areas of high distortion. *Magnetic Resonance Imaging*, 92:1–9, October 2022. ISSN 0730-725X.

Arnaud Boré, Francois Rheault, and Alex Valcourt Caron. Scilpy. Sherbrooke Connectivity Imaging Lab, November 2025.

Michal Byra, Charissa Poon, Muhammad Febrian Rachmadi, Matthias Schlachter, and Henrik Skibbe. Exploring the performance of implicit neural representations for brain image registration. *Sci Rep*, 13(1):17334, October 2023. ISSN 2045-2322.

Leon Y. Cai, Qi Yang, Praitayini Kanakaraj, Vishwesh Nath, Allen T. Newton, Heidi A. Edmonson, Jeffrey Luci, Benjamin N. Conrad, Gavin R. Price, Colin B. Hansen, Cailey I. Kerley, Karthik Ramadass, Fang-Cheng Yeh, Hakmook Kang, Eleftherios Garyfallidis, Maxime Descoteaux, Francois Rheault, Kurt G. Schilling, and Bennett A. Landman. MASiVar: Multisite, multiscanner, and multisubject acquisitions for studying variability in diffusion weighted MRI. *Magnetic Resonance in Medicine*, 86(6):3304–3320, 2021. ISSN 1522-2594.

H. Chang and J.M. Fitzpatrick. A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE Transactions on Medical Imaging*, 11(3): 319–329, September 1992. ISSN 1558-254X.

Matthew Cieslak, Philip A. Cook, Xiaosong He, Fang-Cheng Yeh, Thijs Dhollander, Azeez Adebimpe, Geoffrey K. Aguirre, Danielle S. Bassett, Richard F. Betzel, Josiane Bourque, Laura M. Cabral, Christos Davatzikos, John A. Detre, Eric Earl, Mark A. Elliott, Shreyas Fadnavis, Damien A. Fair, Will Foran, Panagiotis Fotiadis, Eleftherios Garyfallidis, Barry Giesbrecht, Ruben C. Gur, Raquel E. Gur, Max B. Kelz, Anisha Keshavan, Bart S. Larsen, Beatriz Luna, Allyson P. Mackey, Michael P. Milham, Desmond J. Oathes, Anders Perrone, Adam R. Pines, David R. Roalf, Adam Richie-Halford, Ariel Rokem, Valerie J. Sydnor, Tinashe M. Tapera, Ursula A. Tooley, Jean M. Vettel, Jason D. Yeatman, Scott T. Grafton, and Theodore D. Satterthwaite. QSIPrep: An integrative platform for preprocessing and reconstructing diffusion MRI data. *Nat Methods*, 18(7):775–778, July 2021. ISSN 1548-7105.

Oscar Esteban, Christopher J. Markiewicz, Ross W. Blair, Craig A. Moodie, A. Ilkay Isik, Asier Erramuzpe, James D. Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, Hiroyuki Oya, Satrajit S. Ghosh, Jessey Wright, Joke Durnez, Russell A. Poldrack, and Krzysztof J. Gorgolewski. fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nat Methods*, 16(1):111–116, January 2019. ISSN 1548-7105.

Farhad Farzaneh, Stephen J. Riederer, and Norbert J. Pelc. Analysis of T2 limitations and off-resonance effects on spatial resolution and artifacts in echo-planar imaging. *Magnetic Resonance in Medicine*, 14(1):123–139, 1990. ISSN 1522-2594.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1126–1135. PMLR, July 2017.

Levin Fritz, Joost Mulders, Hester Breman, Judith Peters, Matteo Bastiani, Alard Roebroeck, Jesper Andersson, John Ashburner, Nikolaus Weiskopf, and Rainer Goebel. Comparison of EPI distortion correction methods at 3T and 7T. In *OHBM Annual Meeting*, Hamburg, Germany, 2014.

Ali Gholipour, Nasser Kehtarnavaz, Kaundinya Gopinath, Richard Briggs, Michael Devous, and Robert Haley. Distortion Correction via Non-rigid Registration of Functional to Anatomical Magnetic Resonance Brain Images. In *2006 International Conference on Image Processing*, pages 1181–1184, October 2006.

Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, October 2013.

Myla D. Goldman, Shanshan Chen, Robert Motl, Rylan Pearsall, Unsong Oh, and J. Nicholas Brenton. Progression risk stratification with six-minute walk gait speed trajectory in multiple sclerosis. *Front. Neurol.*, 14, October 2023. ISSN 1664-2295.

Mark S. Graham, Ivana Drobnjak, Mark Jenkinson, and Hui Zhang. Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI. *PLOS ONE*, 12(10):e0185647, October 2017. ISSN 1932-6203.

Xuan Gu and Anders Eklund. Evaluation of Six Phase Encoding Based Susceptibility Distortion Correction Methods for Diffusion MRI. *Front. Neuroinform.*, 13, December 2019. ISSN 1662-5196.

Melissa W. Haskell, Jon-Fredrik Nielsen, and Douglas C. Noll. Off-resonance artifact correction for MRI: A review. *NMR in Biomedicine*, 36(5):e4867, 2023. ISSN 1099-1492.

Renaud Hédouin, Olivier Commowick, Elise Bannier, Benoit Scherrer, Maxime Taquet, Simon K. Warfield, and Christian Barillot. Block-Matching Distortion Correction of Echo-Planar Images With Opposite Phase Encoding Directions. *IEEE Transactions on Medical Imaging*, 36(5):1106–1115, May 2017. ISSN 1558-254X.

Dominic Holland, Joshua M. Kuperman, and Anders M. Dale. Efficient correction of inhomogeneous static magnetic field-induced distortion in Echo Planar Imaging. *NeuroImage*, 50(1):175–183, March 2010. ISSN 1053-8119.

Andrew Hoopes, Jocelyn S. Mora, Adrian V. Dalca, Bruce Fischl, and Malte Hoffmann. SynthStrip: Skull-stripping for any brain image. *NeuroImage*, 260:119474, October 2022. ISSN 1053-8119.

Zhangxuan Hu, Yishi Wang, Zhe Zhang, Jieying Zhang, Huimao Zhang, Chunjie Guo, Yuejiao Sun, and Hua Guo. Distortion correction of single-shot EPI enabled by deep-learning. *NeuroImage*, 221:117170, November 2020. ISSN 10538119.

M. Okan Irfanoglu, Pooja Modi, Amritha Nayak, Elizabeth B. Hutchinson, Joelle Sarlls, and Carlo Pierpaoli. DR-BUDDI (Diffeomorphic Registration for Blip-Up blip-Down Diffusion Imaging) method for correcting echo planar imaging distortions. *NeuroImage*, 106:284–299, February 2015. ISSN 1053-8119.

Peter Jezzard and Robert S. Balaban. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*, 34(1):65–73, 1995. ISSN 1522-2594.

Marina Manso Jimeno, Keren Bachi, George Gardner, Yasmin L. Hurd, John Thomas Vaughan Jr., and Sairam Geethanath. GDCNet: Calibrationless geometric distortion correction of echo planar imaging data using deep learning, February 2024.

Heidi Johansen-Berg and Timothy E. J. Behrens. *Diffusion MRI: From Quantitative Measurement to in-Vivo Neuroanatomy*. Elsevier/Academic Press, Amsterdam, 2nd ed edition, 2014. ISBN 978-0-12-396460-1.

J. Kybic, P. Thevenaz, A. Nirkko, and M. Unser. Unwarping of unidirectionally distorted EPI images. *IEEE Transactions on Medical Imaging*, 19(2):80–93, February 2000. ISSN 1558-254X.

Pamela J. LaMontagne, Tammie LS Benzinger, John C. Morris, Sarah Keefe, Russ Hornbeck, Chengjie Xiong, Elizabeth Grant, Jason Hassenstab, Krista Moulder, Andrei G. Vlassenko, Marcus E. Raichle, Carlos Cruchaga, and Daniel Marcus. OASIS-3: Longitudinal Neuroimaging, Clinical, and Cognitive Dataset for Normal Aging and Alzheimer Disease, December 2019. ISSN 1901-4902.

Antoine Legouhy, Mark Graham, Michele Guerreri, Whitney Stee, Thomas Villemonteix, Philippe Peigneux, and Hui Zhang. Correction of Susceptibility Distortion in EPI: A Semi-supervised Approach with Deep Learning. In Suheyla Cetin-Karayumak, Daan Christiaens, Matteo Figini, Pamela Guevara, Tomasz Pieciak, Elizabeth Powell, and Francois Rheault, editors, *Computational Diffusion MRI*, pages 38–49, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-21206-2.

Simin Liu, Yuhui Xiong, Erpeng Dai, Jieying Zhang, and Hua Guo. Improving distortion correction for isotropic high-resolution 3D diffusion MRI by optimizing Jacobian modulation. *Magnetic Resonance in Medicine*, 86(5):2780–2794, 2021. ISSN 1522-2594.

Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, September 2018.

José V. Manjón, Pierrick Coupé, Luis Martí-Bonmatí, D. Louis Collins, and Montserrat Robles. Adaptive non-local means denoising of MR images with spatially varying noise levels. *Journal of Magnetic Resonance Imaging*, 31(1):192–203, 2010. ISSN 1522-2586.

D. Marr and E. Hildreth. Theory of edge detection. *Proc Biol Sci*, 207(1167):187–217, February 1980. ISSN 0080-4649.

Rylan Pearsall, Tyler Spears, Tom Fletcher, Shanshan Chen, Unsong Oh, and Myla Goldman. Validation of a Risk Stratification Tool for MS Progression. In *Multiple Sclerosis Journal*, volume 30, West Palm Beach, Florida, 2024.

Philippe Poulin, Guillaume Theaud, Francois Rheault, Etienne St-Onge, Arnaud Bore, Emmanuelle Renauld, Louis de Beaumont, Samuel Guay, Pierre-Marc Jodoin, and Maxime Descoteaux. TractoInferno: A large-scale, open-source, multi-site database for machine learning dMRI tractography, November 2021.

Yuchuan Qiao and Yonggang Shi. Unsupervised Deep Learning for FOD-Based Susceptibility Distortion Correction in Diffusion MRI. *IEEE Transactions on Medical Imaging*, 41(5):1165–1175, May 2022. ISSN 1558-254X.

Shihan Qiu, Radu Miron, Yahang Li, Cornelius Eichner, Thorsten Feiweier, Nirmal Janardhanan, Bryan Clifford, Omar Darwish, Mahmoud Mostapha, and Mariappan S. Nadar. Unsupervised Learning-Based Susceptibility Artifact Correction for Diffusion-Weighted MRI in Multiple Organs. In James C. Gee, Daniel C. Alexander, Jaesung Hong, Juan Eugenio Iglesias, Carole H. Sudre, Archana Venkataraman, Polina Golland, Jong Hyo Kim, and Jinah Park, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*, pages 635–644, Cham, 2026. Springer Nature Switzerland. ISBN 978-3-032-05169-1.

Paul J. Reber, Eric C. Wong, Richard B. Buxton, and Lawrence R. Frank. Correction of off resonance-related distortion in echo-planar imaging using EPI-based field maps. *Magnetic Resonance in Medicine*, 39(2):328–330, 1998. ISSN 1522-2594.

G.k. Rohde, A.s. Barnett, P.j. Basser, S. Marenco, and C. Pierpaoli. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 51(1):103–114, 2004. ISSN 1522-2594.

Jessica Royer, Raúl Rodríguez-Cruces, Shahin Tavakol, Sara Larivière, Peer Herholz, Qiongling Li, Reinder Vos de Wael, Casey Paquola, Oualid Benkarim, Bo-yong Park, Alexander J. Lowe, Daniel Margulies, Jonathan Smallwood, Andrea Bernasconi, Neda Bernasconi, Birgit Frauscher, and Boris C. Bernhardt. An Open MRI Dataset for Multiscale Neuroscience, August 2021.

D. Rueckert, L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, August 1999. ISSN 1558-254X.

Kurt G. Schilling, Justin Blaber, Yuankai Huo, Allen Newton, Colin Hansen, Vishwesh Nath, Andrea T. Shafer, Owen Williams, Susan M. Resnick, Baxter Rogers, Adam W. Anderson, and Bennett A. Landman. Synthesized b0 for diffusion distortion correction (Synb0-DisCo). *Magnetic Resonance Imaging*, 64:62–70, December 2019. ISSN 0730725X.

Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit Neural Representations with Periodic Activation Functions, June 2020.

Stephen M. Smith. Fast robust automated brain extraction. *Human Brain Mapping*, 17(3): 143–155, 2002. ISSN 1097-0193.

C. Studholme, R.T. Constable, and J.S. Duncan. Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. *IEEE Transactions on Medical Imaging*, 19(11):1115–1127, November 2000. ISSN 1558-254X.

Shanlin Sun, Kun Han, Chenyu You, Hao Tang, Deying Kong, Junayed Naushad, Xiangyi Yan, Haoyu Ma, Pooya Khosravi, James S. Duncan, and Xiaohui Xie. Medical image registration via neural fields. *Medical Image Analysis*, 97:103249, October 2024. ISSN 1361-8415.

Ran Tao, P. Thomas Fletcher, Samuel Gerber, and Ross T. Whitaker. A Variational Image-Based Approach to the Correction of Susceptibility Artifacts in the Alignment of Diffusion Weighted and Structural MRI. In *Information Processing in Medical Imaging*, pages 664–675, Berlin, Heidelberg, 2009. Springer. ISBN 978-3-642-02498-6.

Jeffrey Mark Treiber, Nathan S. White, Tyler Christian Steed, Hauke Bartsch, Dominic Holland, Nikdokht Farid, Carrie R. McDonald, Bob S. Carter, Anders Martin Dale, and Clark C. Chen. Characterization and Correction of Geometric Distortions in 814 Diffusion Weighted Images. *PLOS ONE*, 11(3):e0152472, March 2016. ISSN 1932-6203.

Nicholas J. Tustison, Brian B. Avants, Philip A. Cook, Yuanjie Zheng, Alexander Egan, Paul A. Yushkevich, and James C. Gee. N4ITK: Improved N3 Bias Correction. *IEEE Transactions on Medical Imaging*, 29(6):1310–1320, June 2010. ISSN 1558-254X.

Louis D. van Harten, Jaap Stoker, and Ivana Išgum. Robust Deformable Image Registration Using Cycle-Consistent Implicit Representations. *IEEE Transactions on Medical Imaging*, 43(2):784–793, February 2024. ISSN 1558-254X.

Jelle Veraart, Dmitry S. Novikov, Daan Christiaens, Benjamin Ades-aron, Jan Sijbers, and Els Fieremans. Denoising of diffusion MRI using random matrix theory. *NeuroImage*, 142:394–406, November 2016. ISSN 1053-8119.

Paul Viola and William M. Wells III. Alignment by Maximization of Mutual Information. *International Journal of Computer Vision*, 24(2):137–154, September 1997. ISSN 1573-1405.

Xin Wan, Grant T. Gullberg, Dennis L. Parker, and Gengsheng L. Zeng. Reduction of geometric and intensity distortions in echo-planar imaging using a multireference scan. *Magnetic Resonance in Medicine*, 37(6):932–942, 1997. ISSN 1522-2594.

Sijia Wang, Daniel J. Peterson, J. C. Gatenby, Wenbin Li, Thomas J. Grabowski, and Tara M. Madhyastha. Evaluation of Field Map and Nonlinear Registration Methods for Correction of Susceptibility Artifacts in Diffusion MRI. *Front. Neuroinform.*, 11, February 2017. ISSN 1662-5196.

Jelmer M. Wolterink, Jesse C. Zwienenberg, and Christoph Brune. Implicit Neural Representations for Deformable Image Registration. In *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*, pages 1349–1359. PMLR, December 2022.

Xinyu Ye, Peipei Wang, Sisi Li, Jieying Zhang, Yuan Lian, Yajing Zhang, Jie Lu, and Hua Guo. Simultaneous Super-Resolution and Distortion Correction for Single-shot EPI DWI using Deep Learning, December 2021.

Benjamin Zahneisen, Kathrin Baeumler, Greg Zaharchuk, Dominik Fleischmann, and Michael Zeineh. Deep flow-net for EPI distortion estimation. *NeuroImage*, 217:116886, August 2020. ISSN 10538119.

Abdallah Zaid Alkilani, Tolga Çukur, and Emine Ulku Saritas. FD-Net: An unsupervised deep forward-distortion model for susceptibility artifact correction in EPI. *Magnetic Resonance in Medicine*, 91(1):280–296, 2024. ISSN 1522-2594.

Xi-Nian Zuo, Jeffrey S. Anderson, Pierre Bellec, Rasmus M. Birn, Bharat B. Biswal, Janusch Blautzik, John C. S. Breitner, Randy L. Buckner, Vince D. Calhoun, F. Xavier Castellanos, Antao Chen, Bing Chen, Jiangtao Chen, Xu Chen, Stanley J. Colcombe, William Courtney, R. Cameron Craddock, Adriana Di Martino, Hao-Ming Dong, Xiaolan Fu, Qiyong Gong, Krzysztof J. Gorgolewski, Ying Han, Ye He, Yong He, Erica Ho, Avram Holmes, Xiao-Hui Hou, Jeremy Huckins, Tianzi Jiang, Yi Jiang, William Kelley, Clare

Kelly, Margaret King, Stephen M. LaConte, Janet E. Lainhart, Xu Lei, Hui-Jie Li, Kaiming Li, Kuncheng Li, Qixiang Lin, Dongqiang Liu, Jia Liu, Xun Liu, Yijun Liu, Guangming Lu, Jie Lu, Beatriz Luna, Jing Luo, Daniel Lurie, Ying Mao, Daniel S. Margulies, Andrew R. Mayer, Thomas Meindl, Mary E. Meyerand, Weizhi Nan, Jared A. Nielsen, David O'Connor, David Paulsen, Vivek Prabhakaran, Zhigang Qi, Jiang Qiu, Chunhong Shao, Zarrar Shehzad, Weijun Tang, Arno Villringer, Huiling Wang, Kai Wang, Dongtao Wei, Gao-Xia Wei, Xu-Chu Weng, Xuehai Wu, Ting Xu, Ning Yang, Zhi Yang, Yu-Feng Zang, Lei Zhang, Qinglin Zhang, Zhe Zhang, Zhiqiang Zhang, Ke Zhao, Zonglei Zhen, Yuan Zhou, Xing-Ting Zhu, and Michael P. Milham. An open science resource for establishing reliability and reproducibility in functional connectomics. *Sci Data*, 1(1):140049, December 2014. ISSN 2052-4463.