Phase Correction and Noise-to-Noise Denoising of Diffusion Magnetic Resonance Images using Neural Networks

Jakub Jurek[1][0000−0003−4008−5267], Andrzej Materka[1][0000−0003−0864−1518], Kamil Ludwisiak[2], Agata Majos[3][0000−0002−4141−0191], and Filip Szczepankiewicz[4][0000−0002−5251−587X]

[1] Institute of Electronics, Lodz University of Technology, Aleja Politechniki 10, PL-93590 Lodz, Poland
[2] Department of Diagnostic Imaging, Independent Public Health Care, Central Clinical Hospital, Medical University of Lodz, Pomorska 251, PL-92213 Lodz, Poland
[3] Department of Radiological and Isotopic Diagnosis and Therapy, Medical University of Lodz, Pomorska 251, PL-92213 Lodz, Poland
[4] Medical Radiation Physics, Lund University, Barngatan 4, 22185 Lund, Sweden
jakub.jurek@p.lodz.pl

Abstract. Diffusion magnetic resonance imaging (dMRI) is an important technique used in neuroimaging. It features a relatively low signal-to-noise ratio (SNR) which poses a challenge, especially at stronger diffusion weighting. A common solution to the resulting poor precision is to average signal from multiple identical measurements. Indeed, averaging the magnitude signal is sufficient if the noise is sampled from a distribution with zero mean value. However, at low SNR, the magnitude signal is increased by the rectified noise floor, such that the accuracy can only be maintained if averaging is performed on the complex signal. Averaging of the complex signal is straightforward in the non-diffusion-weighted images, however, in the presence of diffusion encoding gradients, any motion of the tissue will incur a phase shift in the signal which must be corrected prior to averaging. Instead, they are averaged in the modulus image space, which is associated with the effect of Rician bias. Moreover, repeated acquisitions further increase acquisition times which, in turn, exacerbate the challenges of patient motion. In this paper, we propose a method to correct phase variations using a neural network trained on synthetic MR data. Then, we train another network using the Noise2Noise paradigm to denoise real dMRI of the brain. We show that phase correction made Noise2Noise training possible and that the latter improved the denoising quality over averaging modulus domain images.

Keywords: Phase correction · Diffusion magnetic resonance imaging · Denoising · Convolutional neural networks · Transfer learning.

1 Introduction

Diffusion magnetic resonance imaging (dMRI, also diffusion-weighted MRI, DW MRI, DWI) is one of the key tools used both in clinical neurological practice and

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

in neuroscientific research [12]. For example, dMRI is used for visualizing the white matter structures of the brain and can characterize features of the tissue microstructure. The main contrast in dMRI comes from the relation between diffusion encoding strength, characterized by the b-value, and the rate at which water in the tissue is diffusing, i.e., in places where the diffusion is fast, the signal is more attenuated and vice versa. High b-values increase the diffusion-based image contrast and can enable more elaborate analysis models, but the trade-off is a lower signal and thus lower SNR. At some b-value, depending on other imaging parameters including resolution, the SNR becomes prohibitively low. Typically, this problem is solved by repeating the same measurement multiple times and averaging the repetitions (the number of which is referred to as NEX).

Although widely used in practice, averaging has significant limitations. First of all, the noise values converge to their mean value, so averaging should be performed at such a step of image reconstruction where noise has zero-mean. Secondly, averaging reduces the standard error of the mean for the signal intensity proportionally to the square root of NEX, but increases the scan time proportionally to the number of repetitions, so it is inefficient. The maximal NEX number may be limited by the hardware, for example to 32 in the Siemens Avanto 1.5T machine. Since scan time is increased by each repetition, the risk of patient motion and patient discomfort grows.

The condition for zero-mean noise is satisfied for complex-valued images before coil signal combination and before computing modulus images. However, in the presence of diffusion encoding gradients, the position of tissue is encoded in the phase. If the position changes due to movement, vibration, pulsation etc., the phase will be incoherent across the object [11]. This is, repeated images ideally have the same intensity modulus, but different values of the complex intensity components. Thus, images acquired with non-zero diffusion weighting cannot be straightforwardly averaged in the complex domain and are instead averaged in the modulus domain. This results in convergence of intensity values to the non-zero mean and formation of the so called noise floor or Rician bias. This effect is not observed for SNR levels greater than 3, but can bias diffusion measures in most noisy image areas where SNR$<$3 [4].

A summary of MR-related phase correction methods was presented in [11]. The problems associated with phase correction in dMRI were analysed in a recent work by Liu et al. [10].

In this paper, we first propose a neural network, called the Phase Shifter, to equalize phase distribution between pairs of image slices taken from two scan repetitions. The Phase Shifter network is trained using transfer learning, based on synthetic echo planar images (EPI) [6]. We showed in a previous work that denoising can be more effectively tackled using neural networks than by averaging [6]. With phase-corrected images of zero-mean noise at hand, we use the Noise2Noise training paradigm [9] to train a N2N Denoiser network without the necessity for noise-free training targets. For both the phase correction and denoising tasks, we use simple and quick-to-train convolutional neural networks [3].

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

In Section 2, we describe the image simulation process, introduce the Phase Shifter and N2N Denoiser networks and how they were trained. In Section 3, we show the results of training of both networks as well as their application on brain dMRI. The Phase Shifter is compared to polynomial fitting. Section 4 closes the paper with a discussion of the results.

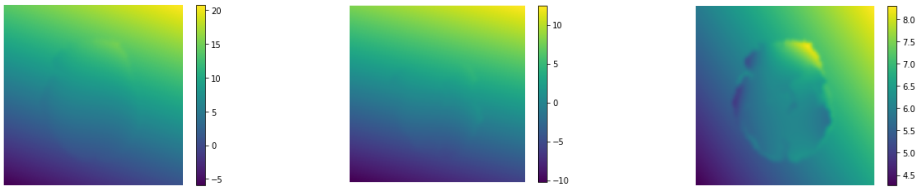

(a) Phase map 1 (b) Phase map 2 (c) Difference of phase maps 1 and 2

Fig. 1: Example simulated phase maps (a, b) and their difference (c), in radians. Phase values can range between several multiples of π. In this work, the goal is to correct image phase by estimating the phase difference, given images with two non-agreeing phase maps (Eq. 1)

2 Materials and methods

Experiments were performed using dMRI of the brain acquired with b-values of 0 and 1000 s/mm^2 on a Siemens Avanto 1.5 T machine at the Central Clinical Hospital, Lodz Medical University, Poland. The scanning involved non-accelerated echo-planar images with imaging parameters described in detail in [6]. Synthetic complex MRI images with random phase were simulated based on the BrainWeb phantom [1][2][7][8] as in [6].

2.1 Modelling synthetic training data for the Phase Shifter

We pose phase correction in dMRI as the problem of finding the phase difference map between two repetitions of the same complex-valued image slice. The ideal, non-noisy modulus value of the complex signal is equal for the two repetitions, up to motion, which we assume negligible in an EPI scan involving two successive acquisitions. The individual real and imaginary values across repetitions are different, but contain noise sampled from a zero-mean distribution.

We performed real-MR-data-driven modelling on the BrainWeb volume to obtain a synthetic training dataset with the required properties [6]. The BrainWeb volume was cropped and downsampled to match the resolution of dMRI data. This resulted in a single 160x160x25 voxels volume. For simplicity in setting the SNR in experiments, the mean value of the tissue in the BrainWeb volume

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

was normalized to 1 using the tissue mask. Phase maps were estimated based on one of the real dMRI scans (Patient 0) by polynomial fitting (PF) [6], cubic within the tissue and linear within background to avoid rapid phase changes in the latter region. In total, 25 phase maps were obtained, one for each slice of the selected dMRI scan. Representative phase maps and their difference are shown in Fig. 1.

The original BrainWeb volume I_m is real valued and thus can be treated as a complex number with zero phase and zero imaginary component. Phase variations within the slices were introduced by:

$$I_c = I_m e^{j\theta} = \mathrm{Re}(I_c) + \mathrm{Im}(I_c)j \tag{1}$$

which yields a complex-valued slices I_c with non-zero real and imaginary parts and phase θ, where θ is a function of pixel coordinates and j is the imaginary unit.

Three different phase maps were imposed on each slice to create triplets and augment the training set (Fig. 2a-2b shows a pair):

$$I_{ci} = I_m e^{j\theta_i}, i = 1, 2, 3 \tag{2}$$

Then, complex noise maps obtained from the scanner were added to the real and imaginary components of all slices. The standard deviation of noise was normalized to 1. EPI regridding effects on noise were taken into account as in [6]. Noisy images with mean SNR = s were calculated as:

$$[I^n_{ci}]^s = s * I_{ci} + N \tag{3}$$

where $[I^n_{ci}]^s$ is the noisy complex image with phase θ_i and SNR s, and N is the zero-mean Gaussian noise. From the triplets of the same slice with different phase, one instance was selected as the one to be corrected and two instances were assumed to be repetitions with reference phase (i=2 and i=1, 3, respectively, by random choice). The goal of $[I^n_{c2}]^s$ phase correction in this setting is to shift its phase from θ_2 to θ_i:

$$([I^n_{c2}]^s)^{\theta_i\prime} = [I^n_{c2}]^s e^{-j(\theta_2 - \theta_i)} \tag{4}$$

Note that $([I^n_{c2}]^s)^{\theta_i\prime}$ has the same phase θ_i but is not equal to $[I^n_{ci}]^s$ due to noise (Fig. 2). By the same argument it satisfies the Noise2Noise criteria for identical image content and zero-mean independent noise. Applying a phase difference computed using the two-argument arctangent, $\mathrm{atan2}([I^n_{c2}]^s) - \mathrm{atan2}([I^n_{ci}]^s)$ (Fig. 3) would yield identical images including their noise values and thus is not suitable. Using two target phase maps for each slice increased the number of training examples to 50. Additionally, SNR values s = 1, 2, 3, 4, 5, 6, 7, 8, 9 were considered, resulting in 450 examples in total. Each input training example was setup as a quadruple $\mathrm{Re}([I^n_{c2}]^s), \mathrm{Im}([I^n_{c2}]^s), \mathrm{Re}([I^n_{ci}]^s), \mathrm{Im}([I^n_{ci}]^s)$, with $s = 1, 2, ..., 9$ and $i = 1, 3$. Targets were setup as tuples $\mathrm{Re}([I^n_{c2}]^s)^{\theta_i\prime}, \mathrm{Im}([I^n_{c2}]^s)^{\theta_i\prime}$.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

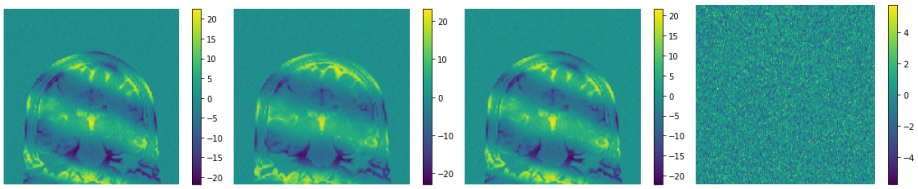

(a) Repetition 1 (b) Repetition 2 (c) (d)

Fig. 2: Two noisy instances of the real part of same slice with different phase (a, b), Repetition 2 with corrected phase (c), the difference of images (a) and (c) which are equal up to noise (d).

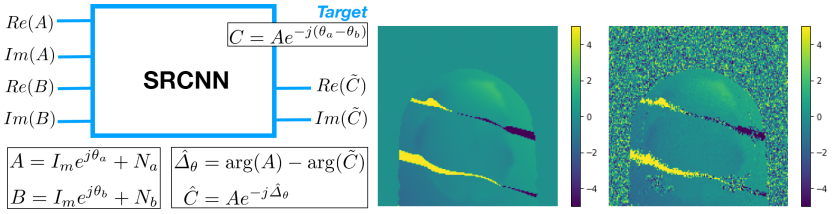

(a) (b) $\theta_a - \theta_b$ (c) $\arg(A) - \arg(B)$

Fig. 3: The scheme of Phase Shifter operation (a) and an example of clean and noisy phase difference maps computed using ideal phase maps (b) and the two-argument arctangent function (c). Using (c) for phase correction would result in two identical image instances, not suitable for averaging nor for Noise2Noise. The desired phase difference map should be noise-free (b)

2.2 Phase correction

The Phase Shifter network architecture selected for correcting the phase is based on the SRCNN [3]. SRCNN is three-layered, with the first layer performing 9x9 convolutions 64 times, the second performing 1x1 convolutions 32 times and the final performing 5x5 convolutions 1 time. The output layer in our Phase Shifter adaptation is modified to perform the convolutions twice, so that it outputs the real and imaginary parts of the phase-corrected image as two channels. The input to the network is a quadruple of images as described above in Section 2.1. (images A and B in Fig. 3a, with $s = 1$). The expected result of the Phase Shifter is a shift in the phase of the first complex input image so that its (unknown) non-noisy complex value is equal to the non-noisy complex value of the second input image.

The Phase Shifter is setup to first output the phase-corrected complex image estimate $\mathrm{Re}((\widehat{[I^n_{c2}]^s})^{\theta_i\prime}), \mathrm{Im}((\widehat{[I^n_{c2}]^s})^{\theta_i\prime})$ (denoted as $\tilde{C}$ in Fig. 3a). Then, the difference is computed between the phase of the input image to be corrected and of the phase-corrected image at the network output, using the two-argument

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

arctangent function:

$$\widehat{(\theta_2 - \theta_i)} = \text{atan2}([I_{c2}^n]^s) - \text{atan2}((\widehat{[I_{c2}^n]^s})^{\theta_i}{}') \tag{5}$$

The latter is applied to the input image to be corrected to obtain the final Phase Shifter output (denoted as $\hat{C}$ in Fig. 3a):

$$(\widehat{[I_{c2}^n]^s})^{\theta_i}{}'' = [I_{c2}^n]^s e^{-j(\widehat{\theta_2 - \theta_i})} \tag{6}$$

The computed phase difference maps $\widehat{(\theta_2 - \theta_i)}$ (denoted as $\hat{\Delta}_\theta$ in Fig. 3a) may initially contain small errors corresponding to discrepancies in the modulus value, but its application keeps the modulus unchanged in the final output, while the errors are minimized as the network is trained.

2.3 Denoising

The Noise2Noise training paradigm is explained in detail in [9]. The main idea is that training a denoiser with clean targets and L2 loss (mean squared error, MSE loss) has the same minimum as training with zero-mean noisy targets. This approach can be applied with any neural network regression-type architecture. For denoising, we again select the SRCNN architecture. We treat the real and imaginary components of a complex image as separate images to denoise. The N2N Denoiser receives a single noisy image slice as input and outputs a single denoised image slice. Its final layer thus performs one 5x5 convolution. The N2N Denoiser estimates the noise map for each image slice and then subtracts it from the input [13]. It is optimized by minimizing the MSE between the denoised images and the noisy phase-corrected targets. The minimum possible loss value is thus, by definition of the MSE, the variance of noise in the target image. Since the images in our experiments were normalized to unit noise standard deviation [6], the optimum loss value is 1.

2.4 Evaluation methods

Phase correction accuracy in synthetic BrainWeb images was evaluated using the Phase Shifter loss function (MSE) globally and by comparing MSE maps in selected cases visually. Phase correction accuracy in dMR images was evaluated by visual examination of the difference between target and phase-corrected images. It was also validated by performing Noise2Noise training using phase-corrected images. Denoising accuracy was evaluated using the MSE calculated in reference to clean images for synthetic BrainWeb images at different SNR levels and to noisy images for real dMRI.

3 Results

3.1 Phase Shifter training and phase correction in dMR images

The Phase Shifter was trained using 450 160x160 synthetic examples using the MSE loss, Adam optimizer, learning rate of 0.001, mini-batch training with batch

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

size of 9 and shuffling, implemented in PyTorch and executed in Google Colaboratory. The validation set was composed of the same image slices, but limited to two phase maps θ_1,θ_2: $\mathrm{Re}([I^n_{c1}]^s), \mathrm{Im}([I^n_{c1}]^s), \mathrm{Re}([I^n_{c2}]^s), \mathrm{Im}([I^n_{c2}]^s)$, and with accordingly modified target: $\mathrm{Re}(([I^n_{c2}]^s)^{\theta_1{}'}), \mathrm{Im}(([I^n_{c2}]^s)^{\theta_1{}'})$.

Thus, there were 225 validation examples. Early stopping was implemented as in [5], with patience of 100 epochs, validation loss value checking in each epoch, and minimum delta value of 1e-9. In this setting, the network reached its best performance after 1645 epochs and training took 44 minutes. Validation mean absolute error (MAE) between the target and estimated phase-corrected images achieved the value of 0.1, which is markedly lower than the standard deviation of the noise (equal to 1). Figs. 4 and 5 show chosen training results.

dMR images of the brain were reconstructed from k-space as in [6], which involved the necessary correction steps: regridding and Nyquist ghosting correction, yielding complex-valued volumes from individual receiver coils. Phase correction was performed using the trained Phase Shifter network on these images (Fig. 6). The network was run 100 times on a pair of 25-slice dMRI volumes and processing duration was approximately 0.02 seconds on average.

Then, polynomial fitting was performed on the same dMRI volumes for comparison. Fitting was run 10 times on each of the volumes. The average run time was approximately 1.55 seconds. Phase distribution after polynomial correction was more similar to the target than prior to correction, but significant errors were noticed.

These errors were quantified by calculating absolute differences between phase-corrected slices and slices with desired phase. Ideally, these differences should be due to noise only. To decrease the influence of noise on this measure and highlight the errors, phase-correction was performed using polynomial fitting and the Phase Shifter on the set of 32 repeated b=1000 s/mm^2 dMR images (a single coil was selected), i.e. 31 volumes were phase corrected and one was selected as reference. Then, the real and imaginary components of all 25 slices and all 32 repetitions were averaged to decrease random noise variance and preserve deterministic content of the error maps (Fig. 7). Averaging these 1600 slice instances is expected to reduce the noise standard deviation by a factor of 40, from 1 to 0.025. The mean absolute error, calculated over the error maps of Fig. 7, was indeed higher for polynomial fitting than it was for the neural network (1.208 vs. 1.127).

3.2 Averaging phase-corrected dMR images

Phase correction was also validated by averaging over the 32 repeated scans. Fig. 8 shows the comparison for the noisy, non-averaged image, averaging without phase correction, averaging after polynomial fitting and after application of the Phase Shifter. For each case, the real and imaginary components are presented as well as the modulus. It can be seen that without phase correction, the intensity of the modulus image greatly deviates from the actual modulus (Figs. 8g and 8c). Both phase correction methods yielded much more similar modulus values. However, it is noticed (Fig. 8l) that polynomial fitting may be locally inaccurate,

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

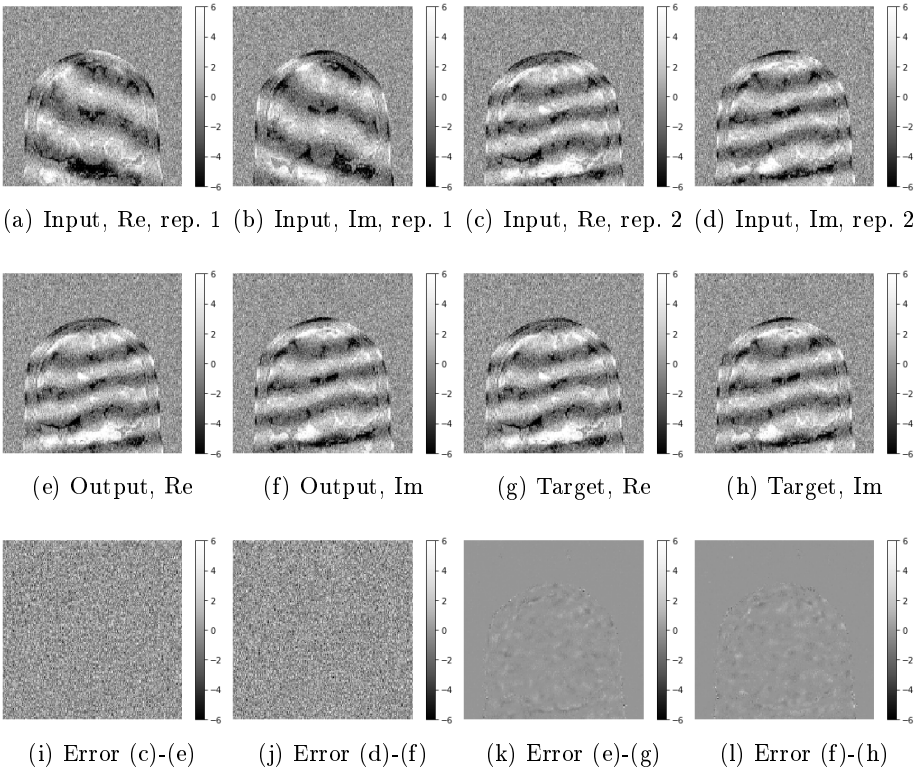

(a) Input, Re, rep. 1 (b) Input, Im, rep. 1 (c) Input, Re, rep. 2 (d) Input, Im, rep. 2

(e) Output, Re (f) Output, Im (g) Target, Re (h) Target, Im

(i) Error (c)-(e) (j) Error (d)-(f) (k) Error (e)-(g) (l) Error (f)-(h)

Fig. 4: Phase shifter training results – a validation example shown at input (a-d) and after Phase Shifter correction compared to targets (e-h). (a-b) is phase-corrected to agree in phase with (c-d). Correction error maps were obtained between the input with target phase and the phase-corrected image (i-j) and between the target and estimated phase-corrected images (k-l). Error maps (i-j) are indistinguishable from noise, which suggests good phase correction. In turn (k-l) show that errors are concentrated within the object and have low value compared to image noise. Re – real part, Im – imaginary part

leading to artifactual intensity variations. It is also visible for both correction methods that they reduced the noise floor compared to the average of slice moduli (Fig. 8d).

3.3 N2N Denoiser training and denoising of brain dMR images

The N2N Denoiser was trained on b=1000 s/mm^2 dMRI data of a single patient, processed in four different ways as described below. dMRI were 160x160x25 in size, acquired for three diffusion directions and using four receiver coils. Each time learning was performed on non-cropped slices, which yielded 300 examples

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

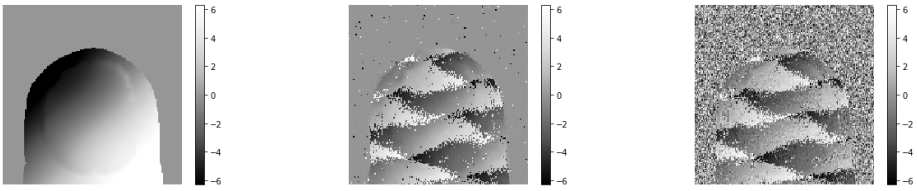

(a) (b) (c)

Fig. 5: Phase difference map target (a), neural network estimate (b) and direct calculation from noisy input images (c). The network seems to learn to estimate a denoised version of the input images, so that Eq. 5 yields less noisy phase estimates (b). Still, both phase estimates exhibit phase wrapping due to the use of arctangent

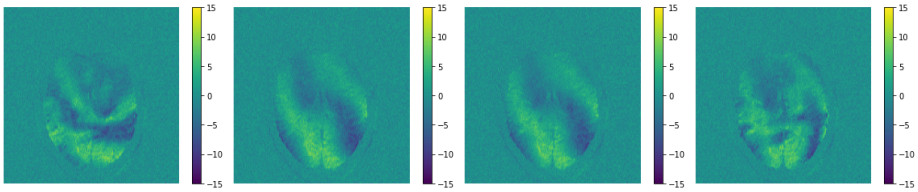

(a) Rep. 1 (b) Rep. 2 (c) Phase shifter, rep. 1 (d) PF, rep. 1

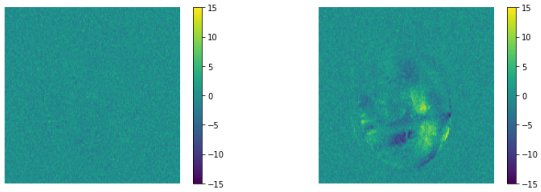

(e) Difference (c)-(b) (f) Difference (d)-(b)

Fig. 6: Example slice from a b=1000 s/mm^2 dMR image, real component. Excellent phase correction is noticed for the Phase Shifter. (e) shows that the correction error is small compared to noise for the Phase Shifter. Errors in phase correction are noticed for polynomial fitting (PF) where it cannot fit the true phase distribution accurately (f).

in total, from which 60 was held out for validation. Hyperparameter settings were MSE loss, Adam optimizer, learning rate of 0.001, mini-batch training with batch size of 9 and shuffling, early stopping with patience of 25 epochs, validation loss value checking in each epoch, and minimum delta value of 1e-6.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

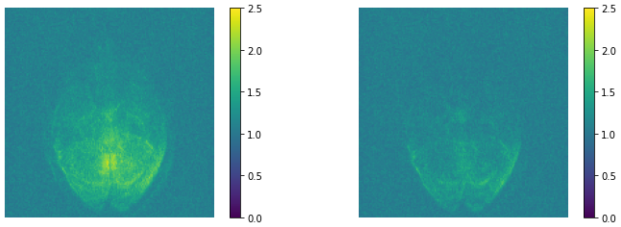

(a) PF, MAE = 1.208

(b) Phase Shifter, MAE = 1.127

Fig. 7: Mean absolute error maps obtained for the b = 1000 s/mm^2 dMRI volume by averaging 1600 images (32 repetitions x 25 slices x 2 complex components)

The N2N denoiser was trained using four kinds of training data: a) without phase correction, b) phase correction with polynomial fitting, c) phase correction with the Phase Shifter, d) no phase correction, moduli of the complex slices were taken as examples.

Additional validation sets were created. They did not have any effect on training and were used to study the N2N learning process with the four kinds of training data. One of them comprised 60 slices of real dMR images, with b = 0 s/mm^2. These dMR images were assumed to have identical phase variation across repetitions. Another four validation sets were synthetic BrainWeb images at SNR = 1, 3, 5 and 7. Validation MSE loss was computed two-fold on these: Noise to Noise loss for noisy targets and Noise to Truth loss (N2T) for clean targets.

Training the network on non-phase-corrected data failed due to substantial differences in the signal phase, which violated Noise2Noise conditions.

Fig. 9 shows the MSE loss between the network output and the noisy phase-corrected image (N2N loss), for the two correction methods, evaluated on the b = 0 s/mm^2 dMRI. Although these images have higher SNR than b = 1000 s/mm^2 data used for training, the network gradually learns to denoise them. The figure shows that the learning process is simpler after phase correction using the Phase Shifter and that the loss is constantly lower than for polynomial fitting. This difference in the loss values may partly result from distortions introduced to the meaningful content of the image by the network trained on images phase-corrected with polynomial fitting. Similar results were obtained for the synthetic validation data at all tested SNR values, where clean images were used as reference (N2T loss, Fig. 10). The fourth network, trained on modulus images, was evaluated by studying the denoised images.

The four versions of the N2N Denoiser were compared based on visual examination of the denoised b = 1000 s/mm^2 images and in comparison to noisy ones. A single-coil and a multicoil image was considered (Fig. 11).

Comparison shows that both phase correction methods significantly improve the SNR of the image compared to the noisy one and to the results of the

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

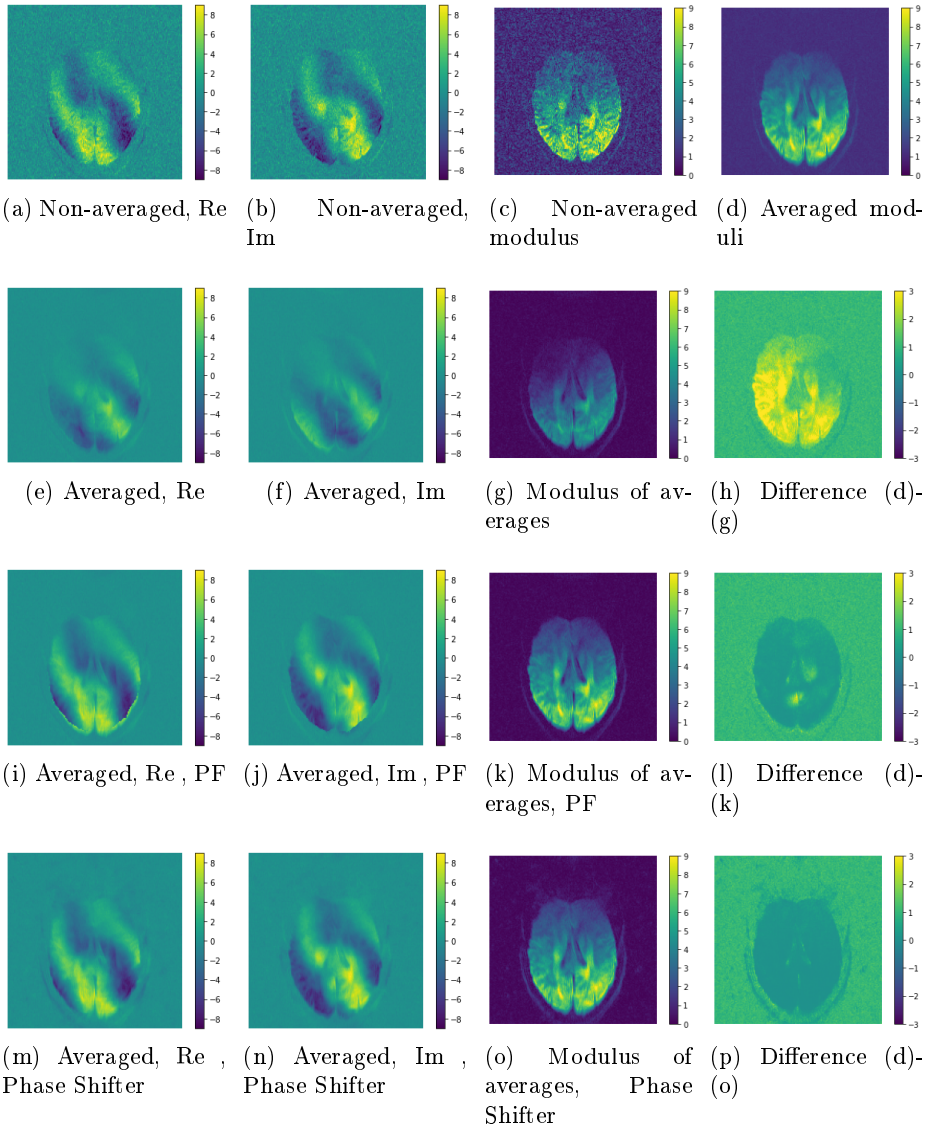

(a) Non-averaged, Re (b) Non-averaged, Im (c) Non-averaged modulus (d) Averaged moduli

(e) Averaged, Re (f) Averaged, Im (g) Modulus of averages (h) Difference (d)-(g)

(i) Averaged, Re , PF (j) Averaged, Im , PF (k) Modulus of averages, PF (l) Difference (d)-(k)

(m) Averaged, Re , Phase Shifter (n) Averaged, Im , Phase Shifter (o) Modulus of averages, Phase Shifter (p) Difference (d)-(o)

Fig. 8: Phase correction results (e-p) after averaging of 32 scan repetitions, compared to a single repetition with desired phase (a-d). Noise floor is reduced due to averaging in the complex domain for both correction methods, but polynomial fitting is not accurate, leading to modulus value variations. Re – real part, Im – imaginary part

N2N denoiser trained on non-phase-corrected data. This is also noticed for the

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

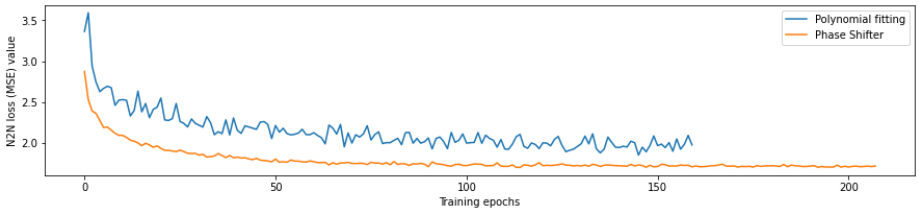

Fig. 9: Validation loss evaluated on the b = 0 s/mm^2 dMRI data while training the N2N denoiser on two kinds of data: corrected with polynomial fitting and the Phase Shifter network

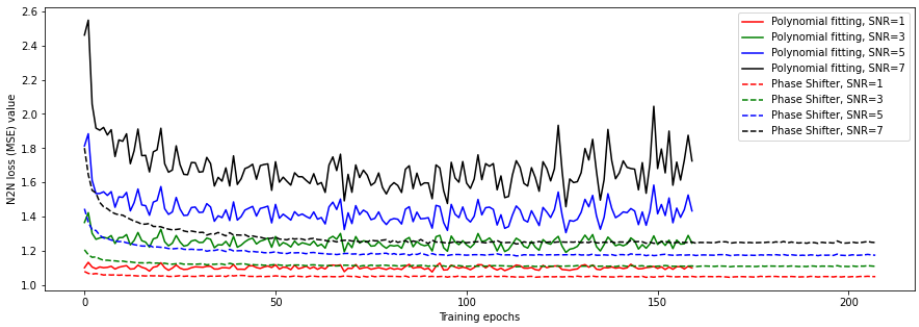

Fig. 10: Validation N2T loss evaluated on the BrainWeb data while training the N2N denoiser on two kinds of data: corrected with polynomial fitting and the Phase Shifter network

network trained on modulus images but Rician bias becomes visible, especially for the multicoil image. Although the N2N Denoiser trained on data that was phase-corrected with polynomial fitting outputs well denoised images, errors in intensity are visible.

4 Discussion

Phase correction is a necessary step to enable complex-domain averaging of dMR images or, as it was shown, Noise to Noise training. Among the two compared methods, polynomial fitting is slower than neural-network-based correction. It requires brain segmentation, because the true background phase is zero and including the background would yield phase discontinuity that is hard to fit accurately using the cubic polynomial. More advanced fitting can be performed, for example using a higher degree polynomial, Legendre polynomials of Radial Basis Functions, which however is even more time-consuming and may involve additional hyperparameters.

The proposed Phase Shifter network was trained on synthetic data, in which phase maps were approximated using polynomial fitting. Despite this, the neural

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

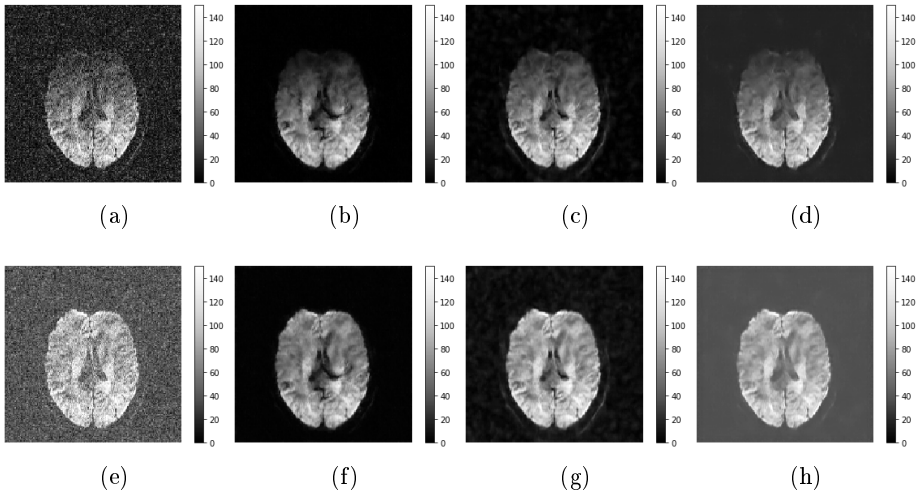

Fig. 11: Noisy image (a, e) and images denoised using the N2N Denoiser trained using different data: with polynomial fitting phase correction (b, f), with Phase Shifter correction (c, g) and with modulus images (d, h). A slice from a selected receiver coil for a b = 1000 s/mm^2 image is shown in the upper row and from a multicoil b = 1000 s/mm^2 image in the lower row

network was able to capture the hidden dependencies and performed better on real data than polynomial fitting, i.e. it was not limited to phase variations that were easily approximated using the polynomial. Fig. 5 shows that the network learns the correct function, i.e. it predicts a denoised phase difference having only noisy complex image pairs as input and a phase-corrected noisy image as target.

Although it may seem more natural to pose the learning problem so that the clean phase maps are the target instead of the phase-corrected noisy images, we encountered significant obstacles using this approach. Namely, the resulting phase maps estimated by the neural network contained non-continuities even if the true phase was continuous, which most probably resulted from phase periodicity. In general, discontinuities of the true phase can also appear naturally due to non-continuous magnetic susceptibility of the tissue, which causes field inhomogeneity. The chosen approach was free of such inaccuracies for continuous true phase and was thus selected. In the future, however, we plan to study in more detail other Phase Shifter training options.

The Phase Shifter uses a two-step approach to estimate the phase-corrected image, in which the first estimate ($\tilde{C}$ in Fig. 3a) is used to compute the phase difference, and then this phase difference is used to compute the final estimate ($\hat{C}$ in Fig. 3a). This solution was applied to ensure that the modulus value of

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

the phase-corrected image was identical to the modulus value of the image to be phase-corrected.

We did not notice phase-correction-related signal loss of the kinde described in [10], in the phase-corrected images obtained using polynomial fitting and the Phase Shifter.

One of the reasons for improving the current network design is other inaccuracies observed from the Phase Shifter phase correction, visible in the results on N2N denoising (Figs. 8, 11). Some background noise was observed to remain in the denoised images, to a greater extent than for polynomial fitting. It is not sure if these distortions are also present within the tissue regions. The probable source is inaccurate estimation of the phase in the background by the Phase Shifter. Although the Phase Shifter network should output zero for the background, it may do so with errors, which may disturb the background noise distribution and lead to noise mean value shifts from zero. In such circumstances, the neural network may learn denoising inaccurately. Alternatively, the network may introduce noise correlation between the N2N inputs and targets. This effect will be more thoroughly studied in our future works. A simple solution would be to use brain masks to exclude the background from phase-shifting, which was the case for polynomial fitting.

Although we tested the phase correction method on b=0 and 1000 s/mm^2 dMRI, it is expected to be applicable, in general, to any b-value or diffusion encoding direction. The limit is probably imposed by decreasing SNR associated with increasing the b-value. Identifying this limit, however, will be the subject of our future studies.

In the future, we also plan to quantify the reduction of bias in the averaged phase-corrected dMRI and the impact of phase correction on image-derived diffusion parameters.

To conclude, neural network phase correction in dMR images is a fast and accurate alternative for polynomial fitting. Training can be performed on simulated images. It enables Noise to Noise denoising in complex dMRI data as well as complex-domain averaging of repeated scans for noise reduction. Future studies are required to study the identified pitfalls and test the phase correction method in more real-life dMRI scenarios.

Acknowledgements Kamil Gorczewski, Kamil Cepuch and Agata Zawadzka from Siemens Healthcare are warmly thanked for their assistance with MR image acquisition and comments on reconstruction.

References

1. Cocosco, C.A., Kollokian, V., Kwan, R.K.S., Pike, G.B., Evans, A.C.: Brainweb: Online interface to a 3D MRI simulated brain database. NeuroImage **5**, 425 (1997)
2. Collins, D., Zijdenbos, A., Kollokian, V., Sled, J., Kabani, N., Holmes, C., Evans, A.: Design and construction of a realistic digital brain phantom. IEEE Transactions on Medical Imaging **17**(3), 463–468 (1998). https://doi.org/10.1109/42.712135

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61

3. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. IEEE Transactions on Pattern Analysis and Machine Intelligence **38**, 295–307 (2016). https://doi.org/10.1109/TPAMI.2015.2439281
4. Gudbjartsson, H., Patz, S.: The rician distribution of noisy MRI data. Magnetic Resonance in Medicine **34**, 910–914 (1995). https://doi.org/10.1002/mrm.1910340618
5. Jurek, J., Kociński, M., Materka, A., Elgalal, M., Majos, A.: CNN-based superresolution reconstruction of 3d MR images using thick-slice scans. Biocybernetics and Biomedical Engineering **40**(1), 111–125 (jan 2020). https://doi.org/10.1016/j.bbe.2019.10.003
6. Jurek, J., Materka, A., Ludwisiak, K., Majos, A., Gorczewski, K., Cepuch, K., Zawadzka, A.: Supervised denoising of diffusion-weighted magnetic resonance images using a convolutional neural network and transfer learning. Biocybernetics and Biomedical Engineering **43**(1), 206–232 (jan 2023). https://doi.org/10.1016/j.bbe.2022.12.006
7. Kwan, R.K.S., Evans, A.C., Pike, G.B.: An extensible MRI simulator for post-processing evaluation. In: Lecture Notes in Computer Science, pp. 135–140. Springer Berlin Heidelberg (1996). https://doi.org/10.1007/bfb0046947
8. Kwan, R.S., Evans, A., Pike, G.: MRI simulation-based evaluation of image-processing and classification methods. IEEE Transactions on Medical Imaging **18**, 1085–1097 (1999). https://doi.org/10.1109/42.816072
9. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: Learning image restoration without clean data. In: Dy, J.G., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018. Proceedings of Machine Learning Research, vol. 80, pp. 2971–2980. PMLR (2018), http://proceedings.mlr.press/v80/lehtinen18a.html
10. Liu, F., Yang, J., Feng, M., Cui, Z., He, X., Zhou, L., Feng, J., Shen, D.: Does perfect filtering really guarantee perfect phase correction for diffusion MRI data? Computerized Medical Imaging and Graphics **103**, 102160 (jan 2023). https://doi.org/10.1016/j.compmedimag.2022.102160
11. Pizzolato, M., Gilbert, G., Thiran, J.P., Descoteaux, M., Deriche, R.: Adaptive phase correction of diffusion-weighted images. NeuroImage **206**, 116274 (feb 2020). https://doi.org/10.1016/j.neuroimage.2019.116274
12. Tax, C.M., Bastiani, M., Veraart, J., Garyfallidis, E., Irfanoglu, M.O.: What's new and what's next in diffusion MRI preprocessing. NeuroImage **249**, 118830 (apr 2022). https://doi.org/10.1016/j.neuroimage.2021.118830
13. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. IEEE Trans. Image Process. **26**(7), 3142–3155 (2017). https://doi.org/10.1109/TIP.2017.2662206

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-36021-3_61