arXiv:2007.13123v1 [eess.IV] 26 Jul 2020

Joint reconstruction and bias field correction for undersampled MR imaging*

Mlanie Gaillochet**, Kerem C. Tezcan**, and Ender Konukoglu

Computer Vision Lab, ETH Zrich, Zrich, Switzerland
gamelani@student.ethz.ch
{tezcan, kender}@vision.ee.ethz.ch

Abstract. Undersampling the k-space in MRI allows saving precious acquisition time, yet results in an ill-posed inversion problem. Recently, many deep learning techniques have been developed, addressing this issue of recovering the fully sampled MR image from the undersampled data. However, these learning based schemes are susceptible to differences between the training data and the image to be reconstructed at test time. One such difference can be attributed to the bias field present in MR images, caused by field inhomogeneities and coil sensitivities. In this work, we address the sensitivity of the reconstruction problem to the bias field and propose to model it explicitly in the reconstruction, in order to decrease this sensitivity. To this end, we use an unsupervised learning based reconstruction algorithm as our basis and combine it with a N4-based bias field estimation method, in a joint optimization scheme. We use the HCP dataset as well as in-house measured images for the evaluations. We show that the proposed method improves the reconstruction quality, both visually and in terms of RMSE.

Keywords: MRI reconstruction · Deep learning · Bias field.

1 Introduction

Magnetic resonance imaging (MRI) is a non-invasive imaging technique that allows studying anatomy and tissue properties without ionizing radiation. However, acquiring a detailed high-quality image is time-consuming.

Reducing scan time is therefore essential in order to increase patient comfort and throughput, and to open up more possibilities for optimized examinations. Since acquisition time in MRI is directly related to the number of samples collected in k-space, shortening it usually means reducing the number of samples collected - for instance, by skipping phase-encoding lines in k-space in the Cartesian case. Yet doing so violates the Nyquist criterion, causing aliasing artifacts and requiring reconstruction techniques that can recover the fully sampled MR image from the undersampled data.

* This work was partially funded by grant 205321_173016 Schweizerischer Nationalfonds zur Frderung der Wissenschaftlichen Forschung

** Equal contribution.

The problem of designing such techniques has received considerable attention in the clinical and signal processing research communities. Conventional approaches involve compressed sensing [3,4] and parallel imaging [5]. More recently, the research efforts have focused on using deep learning to tackle the problem and achieved state-of-the-art results. However, despite their multiple advantages, these learning based algorithms are susceptible to discrepancies between the training set and the image to reconstruct at test time. One such difference is due the different bias fields present in the images. Such discrepancies, sometimes referred to as 'domain shift', typically lead to a loss in performance. Various methods have been developed to tackle the domain shift problem in the deep learning community. One approach is to augment the training data to hopefully resemble the test data better [18][17]. Although this approach increases the generalization capabilities of the learned model, it does not optimize the solution specifically for the test image at hand. Another approach is to modify the network parameters to reduce the domain gap [19], which indirectly optimizes the solution by modifying the network rather than the solution.

Yet, while the effect of the bias field on the image encoding is easy to model and there are readily available methods that can estimate it well, none of the aforementioned approaches takes these into account.

In this work we propose a joint reconstruction algorithm, which estimates and explicitly models the bias field throughout the reconstruction process. By doing so, we remove one degree of variation between the training set and the image to be reconstructed. The training is done on images without bias field, and the bias field itself is modeled as a multiplicative term in the image encoding process, linking the training and test domains. In order to be able to do this we use a reconstruction algorithm which decouples the learned prior information from the image generation process, namely the DDP algorithm [2]. Using the N4 algorithm, we iteratively estimate the bias field in the test image throughout the reconstruction, and this estimation improves as the reconstructed image become better. We compare the proposed method to reconstruction without bias field estimation, using a publicly available dataset as well as in-house measured images, and show improvement in performance.

2 Methods

2.1 Problem formulation

The measured k-space data, $\mathbf{y} \in \mathcal{C}^M$ and the underlying true MR image, $\mathbf{x} \in \mathcal{C}^N$ are related through the encoding operation $\mathbf{E} : \mathcal{C}^N \to \mathcal{C}^M$ (which incorporates the coil sensitivites, Fourier transformation and the undersampling operation) as $\mathbf{y} = \mathbf{E}\mathbf{x} + \eta$, where η is complex Gaussian noise.

2.2 DDP reconstruction

The deep density prior reconstruction is based on maximum-a-posteriori estimation, i.e. $\max_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y}) = \max_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x})$, where $\log p(y|x)$ is the

data consistency and $\log p(x)$ is the prior term. The data consistency can be written exactly as the log Gaussian and the method approximates the prior term using the evidence lower bound (ELBO) of a variational auto-encoder (VAE) [24][23]. With these, the reconstruction problem becomes

$$\min_{x} ||\mathbf{E}\mathbf{x} - \mathbf{y}|| - ELBO(\mathbf{x}). \tag{1}$$

The VAE is trained on patches from fully sampled images and the ELBO term operates on a set of overlapping patches that cover the whole image. As evident in the equation above, the ELBO term is independent of the image encoding. The DDP method solves the problem using the projection onto convex sets (POCS) algorithm. In this scheme the optimization is implemented as successive applications of projection operations for prior $\mathcal{P}_{prior}$, data consistency $\mathcal{P}_{DC}$ and the phase of the image $\mathcal{P}_{phase}$, i.e. $\mathbf{x}^{t+1} = \mathcal{P}_{DC}\mathcal{P}_{phase}\mathcal{P}_{prior}\mathbf{x}^t$. The prior projection is defined as a gradient ascent for a fixed number of steps for the ELBO term w.r.t. the image magnitude, i.e. $\mathcal{P}_{prior}\mathbf{x} = \mathbf{x}^N$, where $\mathbf{x}^{n+1} = \mathbf{x}^n + \alpha \frac{d}{d\mathbf{x}} ELBO(|\mathbf{x}|)|_{\mathbf{x}=\mathbf{x}^n}$ for $n = 0...N$. The data consistency projection is given as $\mathcal{P}_{DC}\mathbf{x} = \mathbf{x} - \mathbf{E}^H(\mathbf{E}\mathbf{x} - \mathbf{y})$. For the phase projection, we use the one defined by [2], to which we add the minimization of the data consistency with respect to the image phase. For further details, we refer the reader to the original paper [2].

2.3 Modeling and estimating the bias field

Signal intensities are often not constant across the MR images, even inside the same tissue. Instead, they usually vary smoothly, with fluctuations of 10%-20%, across the measured 3D volume [6]. These variations, collectively called the bias field, can be attributed to factors like patient anatomy, differences in coil sensitivity or standing wave effects [6,9], which are difficult or impossible to control during an acquisition. Hence they can introduce a degree of variation between images used for training reconstruction models and an image to be reconstructed at test time. In order to prevent a loss of performance, this variation has to be taken into account.

To this end, we model the bias field explicitly and incorporate it into Equation 1 as a multiplicative term $\mathbf{B}$ before the encoding operation, modifying the image intensities pixelwise. The bias field is an additional unknown in the reconstruction process that is estimated alongside x as

$$\min_{\mathbf{x},\mathbf{B}} ||\mathbf{E}\mathbf{B}\mathbf{x} - \mathbf{y}|| - ELBO(\mathbf{x}). \tag{2}$$

Here, the measured k-space $\mathbf{y}$ carries the effect of a bias field, while $\mathbf{x}$ is bias field free as the bias field in the image is explicitly modeled using the $\mathbf{B}$ term. This setting allows us to learn the VAE model on images without bias field. The advantage of this idea is two fold. Firstly, since training can be performed on bias field free images, it is easier for the VAE to learn the distribution as there is less spurious variation in the data. Secondly, we make the reconstruction problem

easier by explicitly providing the bias field information, which otherwise would have to be reconstructed from the undersampled k-space as well. We solve Equation 2 as a joint iterative reconstruction problem by minimizing alternatively two sub-problems:

$$1. \quad \mathbf{x}^t = \min_x ||\mathbf{E}\mathbf{B}^{t-1}\mathbf{x} - \mathbf{y}|| - ELBO(\mathbf{x}) \tag{3}$$

$$2. \quad \mathbf{B}^t = \mathrm{N4}(\mathbf{B}^{t-1}\mathbf{x}^t), \tag{4}$$

where N4 denotes the bias field estimation algorithm, which we will explain below. To account for the bias field, the data consistency projection $\mathcal{P}_{DC}$ needs to be adapted and becomes $\mathcal{P}^B_{DC}\mathbf{x} = \mathbf{B}^{-1}[\mathbf{B}\mathbf{x} - \mathbf{E}^H(\mathbf{E}\mathbf{B}\mathbf{x} - \mathbf{y})]$. In this case the reconstructed image corresponding to $\mathbf{y}$ is given as $\mathbf{B}\mathbf{x}$. This modification can be interpreted as doing a forward-backward projection with $\mathcal{P}_{bias} = \mathbf{B}$ before and after the data consistency projection to move the image between the "normalized" bias field free domain and the bias field corrupted acquisition domain, i.e. $\mathbf{x}^{t+1} = \mathcal{P}^{-1}_{bias}\mathcal{P}_{DC}\mathcal{P}_{bias}\mathcal{P}_{phase}\mathcal{P}_{prior}\mathbf{x}^t$. The pseudocode for the described joint optimization scheme is presented in Algorithm 1.

N4 bias field estimation N4 is a variant of the widely used N3 algorithm [7], a non-parametric iterative method that approximates intensity non-uniformity fields in 3-dimensional images. Given an image, its objective is to find a smooth, slowly varying, multiplicative field [6]. N4 improves upon the N3 algorithm by modifying the B-spline smoothing strategy and the iterative optimization scheme used in the original framework. We use the implementation available as N4ITK [8]. We denote this method as $N4(\cdot)$ in our formulations.

Algorithm 1 Joint reconstruction

```
y ← undersampled k-space data
E← undersampling encoding operator
VAE ← trained VAE
NumIter, BiasEstimFreq, DCProjFreq
procedure JointRecon(y, E, VAE)
    B ← N4(E^H y)
    x^0 ← (B)^{-1} E^H y
    for t: 0 to NumIter - 1 do
        x^{t+1} = P_prior x^t
        x^{t+1} ← P_phase x^{t+1}                                  ▷ Optional
        if t% DCProjFreq == 0 and t ≠ 0 then
            x^{t+1} ← B^{-1} [B x^{t+1} − E^H (E B x^{t+1} − y)]   ▷ P^B_DC
        if t% BiasEstimFreq == 0 and t ≠ 0 then
            B ← N4(B x^{t+1})
    return x^{t+1}, B
```

2.4 Datasets used

To train the VAE we used 5 central, non-adjacent slices with 0.7mm isotropic resolution from T1 weighted images of 360 subjects from the Human Connectome Project (HCP) preprocessed dataset [15,16], which by default have a bias field. We used N4 on the images to also create a bias field free training set.

For test images, we took a central slice from 20 different test subjects from the HCP data. As the HCP images tend to have a similar bias field, we additionally created a modified test set where we estimated the bias fields with N4, took their inverse and multiplied them with the bias field free images. In addition to HCP data, we also tested the proposed method with central slices from 9 in-house measured subjects. These images were acquired using a 16 element head coil and have similar acquisition parameters as the HCP dataset with a 1mm isotropic resolution. We used ESPIRiT [21] to estimate the coil sensitivity maps for these images.

2.5 Training the VAE

We trained four patch-wise VAEs - for two different resolution levels to match the datasets, each with and without bias field. We used patches of size 28x28 with a batch size of 50 and ran the training for 500,000 iterations. The patches were extracted randomly from the training images with replacement. All the VAEs were trained with the same training images extracted from the HCP dataset, as described above.

2.6 Experimental Setup

We used random Cartesian undersampling patterns with 15 fully sampled central profiles. We generated a different pattern for each subject, and applied the same pattern for a given subject throughout all experiments for comparability of results. When reconstructing test images from the HCP dataset, we used 302 iterations (NumIter) for R=2, 602 for R=3, 1002 for R=4 and 1502 for R=5, to allow for convergence of reconstructed images. Since the in-house measured images have multiple coils, the successive applications of data consistency projections coincide with a POCS-SENSE reconstruction [20], speeding up convergence and requiring less iterations. Hence, when performing reconstruction on images from the in-house measured dataset, we ran the first 10 iterations without prior projection, applying only data consistency projections. Additionally, the discrepancy between the actual coil sensitivities and the ESPIRiT [21] estimations may lead to divergence after too many iterations. Hence, a lower number of iterations was used for reconstruction experiments on the in-house measure dataset: 32 iters for R=2, 102 for R=3 and R=4, and 202 for R=5. For all test datasets, the parameters were set as α = 1e-4, BiasEstimFreq=10, DCProjFreq=10. As for the N4 bias field estimation algorithm, the default parameters were used.

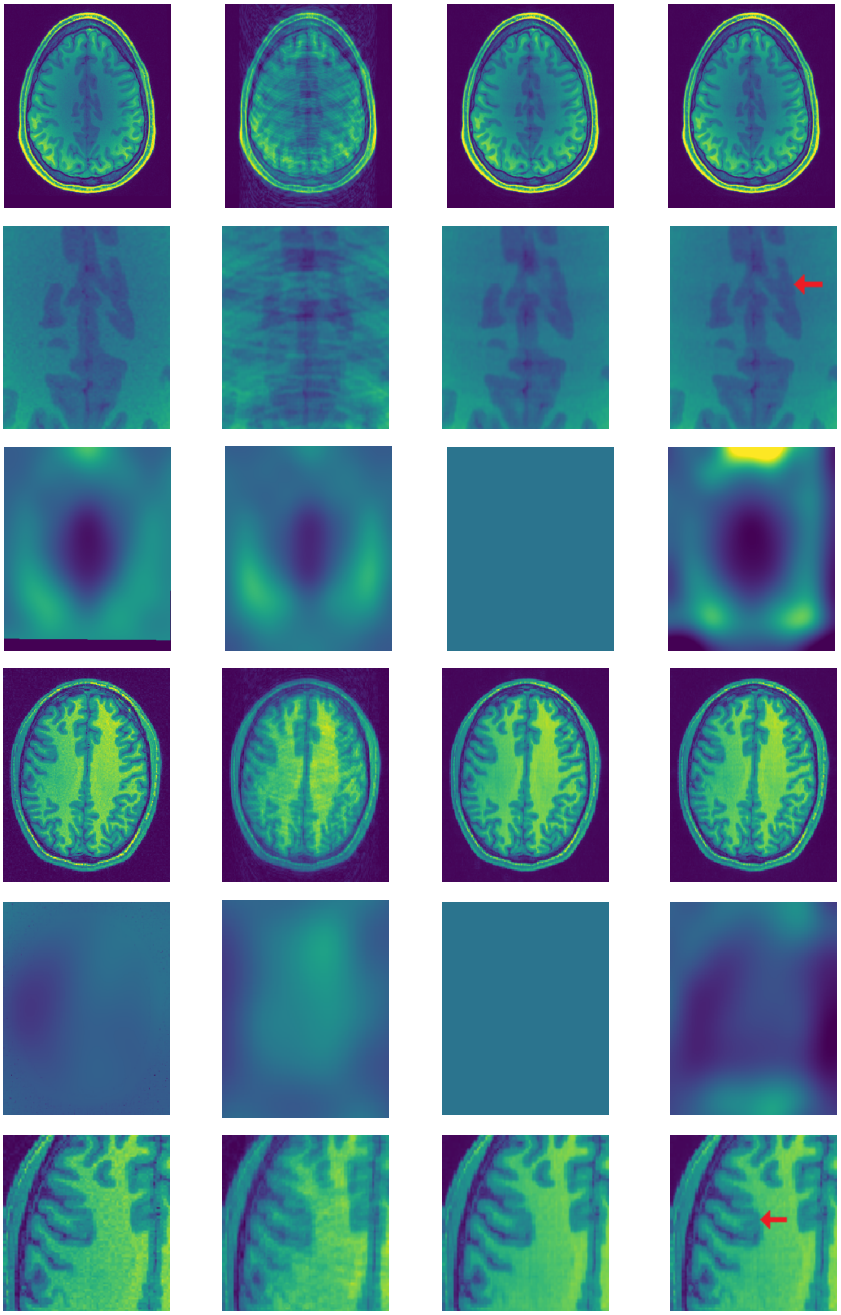

(a) FS (b) Zero-filled (c) Baseline (d) Joint recon.

Fig. 1: Reconstruction results for R = 3, with (a) the fully sampled image, (b) the zero-filled image, (c) the reconstruction with no bias field estimation, (d) the joint reconstruction with bias field estimation using N4. The first three rows show reconstruction results for an HCP image, its zoomed-in version and the corresponding bias field. The next three rows show results for an in-house measured image. For visualization purposes, MR images are clipped to [0, 1.2] and bias fields, to [0.5, 1.8].

To evaluate the reconstruction quality of the different models, we computed, for each image, the percentage RMSE. More specifically, the evaluation metric was given as $\text{RMSE} = 100 \times \sqrt{\frac{\sum_i (|\mathbf{Bx}| - |\hat{\mathbf{x}}|)_i^2}{\sum_i |\hat{\mathbf{x}}|_i^2}}$ where $\hat{\mathbf{x}}$ is the original, fully sampled test image, and where the summation was applied pixel-wise. The error was computed on the skull-stripped images only in the brain area.

To evaluate the statistical significance of our results, we performed a permutation test [22] with 10,000 permutations to assess the null hypothesis that the RMSE's for the reconstruction with and without bias field estimation are from the same distribution. From these tests, we reported the p-values.

3 Experiments and Results

To assess the hypothesis that correcting the bias field improves the overall reconstruction quality, we performed reconstructions on the original and modified HCP test sets as well as on the in-house measured images.

When applying our proposed method, we reconstructed undersampled images with a bias field, and used a VAE trained on images without bias field as well as the N4 bias field estimation algorithm. To evaluate our approach, we also ran baseline experiments on the same images, where we used a VAE trained on images with a bias field and did not apply bias field estimation during reconstruction. Given that the ground truth (fully sampled) images naturally have a bias field in them, we utilized the last bias field estimate multiplied with the reconstructed image, i.e. $\mathbf{Bx}$, for visualisations and RMSE calculations.

Table 1: Table of RMSE values. R is the undersampling factor. Numbers indicate the mean (std). The * indicates a p-value of less than 0.05. The baseline method is the DDP algorithm described in [2] that does not explicitly model the bias field. The proposed joint reconstruction method estimates the bias field using N4 and explicitly models it during reconstruction.

Method	HCP dataset				Modified HCP dataset				In-house measured dataset			
	R= 2	R = 3	R = 4	R = 5	R= 2	R = 3	R = 4	R = 5	R= 2	R = 3	R = 4	R = 5
Baseline	2.24 (0.31)	3.39 (0.45)	4.42 (0.51)	5.72 (1.05)	2.24 (0.38)	3.45 (0.60)	4.26 (0.46)	5.58 (1.20)	4.64 (0.391)	6.852 (0.726)	8.593 (1.344)	11.046 (1.727)
Joint recon.	2.27 (0.34)	3.34 (0.40)	4.35* (0.47)	5.52 (0.66)	2.20* (0.39)	3.33* (0.53)	4.16 (0.54)	5.12* (0.57)	4.62 (0.418)	6.714* (0.821)	8.218* (1.266)	10.567* (1.632)

The results in Table 1 indicate that the proposed joint reconstruction method with bias field estimation improves the reconstruction quality in terms of RMSE when the bias field of the test image is different from those in the training set. In these cases, namely the experiments with the modified HCP and in-house measured images, the improvement is statistically significant with a p-value of less than 0.05 for nearly all undersampling factors. For the unmodified HCP dataset, where the bias field in the test images matches those in the training set,

we do not expect a big difference in the performance, which is reflected in the results.

The quantitative improvement is also supported by the visual inspection of the images given in Figure 1. From the HCP image, one can observe that the level of artifacts is reduced with the proposed method. This becomes more evident in the zoomed-in images. The red arrow points to a part of the image where the proposed method can reconstruct the structures faithfully, whereas the baseline method struggles. Aliasing artifacts are globally suppressed better with the joint reconstruction method. Similarly, for the in-house measured image, the grey matter structure that the red arrow points to is not reconstructed in the baseline method, whereas it again appears with the proposed method.

In this work, we demonstrate the performance loss due to the bias field for a specific algorithm, into which we also integrate our proposed solution. However, the problem may not be specific to the algorithm used as it arises from the domain gap, which is a fundamental problem affecting machine learning based methods in general. Furthermore, the proposed method of estimating and explicitly modeling the bias field in reconstruction is also a generic approach, which can be integrated into different algorithms.

4 Conclusion

In this paper we proposed and evaluated a method for joint reconstruction and bias field estimation to address variations due to the bias field when reconstructing undersampled MR images. The results indicate that the proposed method improves the baseline method (unsupervised learning based reconstruction algorithm), both in RMSE and visually. The improvements can be attributed to two factors. First, the proposed method allows the VAE prior to learn a simpler distribution, and second, providing the bias field explicitly makes the reconstruction problem easier. In essence, estimating the bias field during reconstruction makes the model less sensitive to differences between the data used to train the model and the test data used during reconstruction.

References

1. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A.W.M. , van Ginneken, B., Sanchez, C. I.: A Survey on Deep Learning in Medical Image Analysis. Medical Image Analysis **42**, 60-88 (2017)
2. Tezcan, K. C., Baumgartner, C. F., Luechinger, R., Pruessmann, K. P., Konukoglu, E.: MR Image Reconstruction Using Deep Density Priors. IEEE Transactions on Medical Imaging **38**(7), 1633–1642 (2019)
3. Donoho, D. L.: Compressed sensing: IEEE Transactions on Information Theory **52**(4), 1289-1306 (2006)
4. Lustig, M., Donoho, D., Pauly, J. M.: Sparse MRI: The application of compressed sensing for rapid MR imaging. Magnetic Resonance in Medicine **58**(6), 1182-1195 (2007)

5. Deshmane, A., Gulani, V., Griswold, M. A., Seiberlich, N.: Parallel MR Imaging. Journal of Magnetic Resonance Imaging **36**(1), 55–72 (2012)
6. Sled, J. G., Zijdenbos, A. P., Evans, A. C.: A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Transactions on Medical Imaging **17**(1), 87–97 (1998)
7. Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., Gee J. C.: N4ITK: Improved N3 Bias Correction. IEEE Transactions on Medical Imaging **29**(6), 1310-1320 (2010)
8. SimpleITK: N4 Bias Field Correction, https://simpleitk.readthedocs.io/en/master/link_N4BiasFieldCorrection_docs.html. Last accessed 12 March 2020
9. Sled J.G., Pike G.B.: Understanding intensity non-uniformity in MRI. In: Wells W.M., Colchester A., Delp S. (eds.) MICCAI 1998, LNCS, vol. 1496, pp. 614–622. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0056247
10. Meyer, C. R., Bland, P., H., Pipe, J.: Retrospective correction of intensity inhomogeneities in MRI. IEEE Transactions on Medical Imaging, **14**(1) (1995)
11. Automated Model-Based Bias Field Correction of MR Images of the Brain. IEEE Transactions on Medical Imaging **18**(10), 885–896 (1999)
12. Leemput, K., V., Maes, F., Vandermeulen, D., Suetens P.: Automated Model-Based Tissue Classification of MR Images of the Brain. IEEE Transactions on Medical Imaging **18**(10), 897–908 (1999)
13. Ashburner, J., Friston, K. J.: Unified segmentation. NeuroImage, **26**(3), 839–851 (2005)
14. The Minimal Preprocessing Pipelines for the HumanConnectome Project. Neuroimage **80**, 105-124 (2013)
15. Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: An Overview. Neuroimage **80**, 62-79 (2013)
16. Human Connectome Reference Manual, https://www.humanconnectome.org/storage/app/media/documentation/s1200/HCP_S1200_Release_Reference_Manual.pdf. Last accessed 4 March 2020
17. Chaitanya K., Karani N., Baumgartner C.F., Becker A., Donati O., Konukoglu E.: Semi-supervised and Task-Driven Data Augmentation. In: Chung A., Gee J., Yushkevich P., Bao S. (eds.) IPMI 2019, LNCS, vol. 11492, pp. 29–41. Springer, Cham (2019)
18. Volpi, R., Namkoong, H., Sener O., Duchi J., Murino V., Savarese, S.: Generalizing to Unseen Domains via Adversarial Data Augmentation. NeurIPS **31**, 5334–5344 (2018)
19. Sun B., Saenko K.: Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In: Hua G., Jgou H. (eds.) ECCV 2016, LNCS, vol. 9915. Springer, Cham (2016)
20. Samsonov AA, Kholmovski EG, Parker DL, Johnson CR: POCSENSE: POCS-based reconstruction for sensitivity encoded magnetic resonance imaging. MRM **52**(6), 1397-1406 (2004)
21. Uecker M., Lai P., Murphy M. J., Virtue P, Elad M, Pauly J. M., Vasanawala S. S., Lustig M: ESPIRiT–an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. MRM **71**(3), 990-1001 (2014)
22. Collingridge, D.S.: A Primer on Quantitized Data Analysis and Permutation Testing. Journal of Mixed Methods Research **7**(1) (2013)
23. Kingma D. P., Welling, M.: Auto-Encoding Variational Bayes. arXiv:1312.6114v10, (2014)

24. Rezende, D. J., Mohamed S., Wierstra D.: Stochastic Backpropagation and Approximate Inference in Deep Generative Models. arXiv:1401.4082v3, (2014)