How to train a Deep Convolutional Neural Network for Quantitative Susceptibility Mapping (QSM)

Thomas Jochmann [1], Jens Haueisen [1], and Ferdinand Schweser [2,3]

[1] Technische Universität Ilmenau, Department of Computer Science and Automation, Ilmenau, Germany; [2] Center for Biomedical Imaging at the Clinical and Translational Science Institute, University at Buffalo, The State University of New York, Buffalo, New York, USA; [3] Buffalo Neuroimaging Analysis Center, Department of Neurology, Jacobs School of Medicine and Biomedical Sciences at the University at Buffalo, The State University of New York, Buffalo, New York, USA

Summary

- Deep neural networks are increasingly being used for QSM dipole inversion, mostly using convolutional architectures such as the *U-Net*.
- Training samples of susceptibility and frequency distributions can be obtained from physical simulations on synthetic source distributions, or through "classical" QSM methods on real data.
- Neural networks mimic the structures found in the training examples and generalize to similar cases. → It is important to train with a large variety of patterns (susceptibility distributions).
- Augmentation increases the variety presented to the network and improves the training effect per iteration epoch. → Higher accuracy, faster convergence.
- Artificial *Rician noise* during training improves the robustness against real-world phase noise.
- Deep neural networks remain difficult to interpret. Synthetic samples reliably protect against anatomical biases, i.e. "learning how to *paint* a healthy brain".

Introduction

Deep convolutional neural networks, especially the *U-Net* architecture [1], are increasingly being used to solve the QSM dipole inversion problem [2, 3].

For QSM, the neural network is trained to approximate the inverse mapping from *magnetic susceptibility* χ to *frequency distribution* f (Fig. 1).

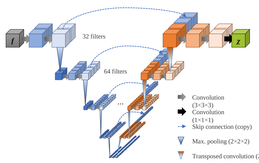

Figure 1. **The U-Net architecture used for QSM dipole inversion.** The network takes frequency distributions as input and outputs magnetic susceptibility distributions.

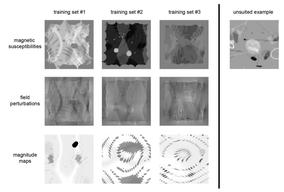

Figure 2. **Synthetic samples.** Left: three examples of synthetic training data sets. From top to bottom: Magnetic susceptibility source distribution, simulated field perturbation, and magnitude images used for creation of Rician (magnitude-modulated) noise. Right: example of a training data variant with too many zero-valued regions in the volume, leading to artifacts during application on realistic brain anatomies

The training examples of χ can be

- synthetic (e.g. random geometric patterns) [2] **or**
- realistic (e.g. from COSMOS) [3]

whereas the corresponding frequency distributions are most commonly obtained via *Fast Forward Field Computation* [4], potentially with added non-susceptibility terms [5, 6].

$$f_{total} = d * \chi \ (+ f_{non\text{-}susceptibility}\ [\text{optionally}])$$

Although the physical fidelity of the learned mapping $f \rightarrow \chi$ is strongly determined by the training examples, practical guidance on hyperparameters, training data simulation and augmentation is scarce.

Here we report on an improved training strategy and present common pitfalls for deep-learning-based QSM:

- *Physics-aware augmentation* to improve the optimization per sample and per iteration epoch, and to prevent overfitting.
- *Artificial noise* to increase the robustness against real noise and to prevent overfitting.
- *Synthetic geometries* to yield infinitely many training samples and to prevent anatomical bias.
- *Theoretical considerations and requirements* regarding sample size and synthetic patterns.
- *Practical experience* on edge treatment and padding, cost function, and hyperparameters.

Methods

NEURAL NETWORK ARCHITECTURE

- *U-Net* [1] *with 3D in- and output* (Tensorflow 2.0.0) to predict the volumetric scalar field χ from a same-sized volumetric scalar field f (Fig. 1).

DATA FOR PHYSICS-INFORMED TRAINING

- *1200 synthetic, entirely non-anatomical samples of* χ (144×144×144 voxels), consisting of randomly shaped patterns with random amplitudes (Fig. 2). Gaussian blurring with random kernel sizes was applied individually to each structure in the volume.
- *Random source free area* ($\chi = 0$) *of approx. 10 pixels was added to the edges* of the volume to ensure that the field decays to zero and therewith *decrease physical errors introduced by zero-padding* (Fig 2).
- f *was obtained by computing the field perturbation* via Fast Forward Field Computation (i.e. dipole convolution in k-space) [4].

AUGMENTATION AND OPTIMIZATION

- We *trained for 500 epochs with ADAM*.
- *Augmented the training data* (Fig. 3) by randomly
  - rotating (90° steps along B_0 axis),
  - flipping (x, y, z),
  - scaling (factors 0.1 to 2), and
  - Rician noise [7].
- *We introduce a custom loss function* that is invariant against amplitude scaling, in order to learn the same from down-scaled and up-scaled examples:

$$L = \| f_{ground\ truth} - f_{prediction} \|^2 / \mathrm{var}(f_{ground\ truth})$$

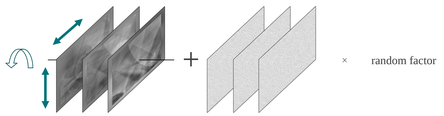

Figure 3. **Augmentation of the training data and added noise.** Augmentation increases the effective training data size and improves final accuracy as well as convergence speed. Each sample of χ and the corresponding f are randomly rotated (along B_0), flipped (x,y,z) and scaled before artificial *Rician noise* is added.

EVALUATION AND APPLICATION

- 120 held-out synthetic samples, similar to the training data, but without augmentation.
- We measured the mean squared error and evaluated visually.
- Datasets from the 2019 QSM Reconstruction Challenge.

Results

- The reconstructed synthetic evaluation patterns were almost indistinguishable from the ground truth (Fig. 4A vs. 4C). The neural network's reconstruction from the challenge dataset (Fig. 5B) was *on par* with state of the art conventional QSM methods.
- Our changes to the training data had a much stronger impact to the accuracy of the learned mapping than tuning the neural network's hyperparameters (kernel counts, activation function, average pooling vs. of max pooling).
- Setting up the U-Net based on 32 kernels in the first layer performed notably better (i.e. lower MSE) than with only 16 kernels, but also better than with 48 or 64 kernels.
- Different representations of the dipole (declaration in the spatial and Fourier domain) had no appreciable impact.
- It was important to have diverse training patterns to cover a wide range of local scenarios, meaning susceptibility gradients and edges in any size and direction.
- Training with rather small-sized structures that are well contained within the zero-background volume (Fig. 6) resulted in poor reconstruction of large white matter regions (Fig. 5D).
- With non-zero susceptibilities up to the boundary of the volume, zero-padding causes physically incorrect training data which lead to artifacts (Fig 5E).
- Adding random noise to the training data reduced the problem of noise amplification.
- Augmenting the training samples strongly improved the mapping as measured by the mean squared error on the reconstruction of the synthetic evaluation set.

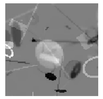

Figure 6. **Synthetic susceptibility distribution with too small structures.**

Discussion

- We expected 16 kernels in the first layer to be too few, as detecting edges in 45-degree steps in three dimensions already requires 26 kernels when using a rectifying activation function. Further kernels are needed to sense average values and fine, pixel-sized structures. Indeed, we found 32 kernels to perform better, but were not able to make use of the potentially higher capacity of 48 or 64 kernels.
- The minimum possible sample size in our U-Net would be 48×48×48, but in that case, the images in the lowest layer would be as small as 3×3×3 = 27 voxels with 98 zero-padded voxels, leaving it hard to train. We used 144×144×144.
- In real data, the field values within calcifications and outside the brain volume do not adhere to the dipole convolution model. These cases are yet to be accounted for in the training samples.

Conclusion

- We presented a simulation and training strategy with physics-aware augmentation and a custom loss function to train a deep convolutional neural network for QSM.
- Improving the physical fidelity of the overall method might require further steps beyond a single training of a neural network on examples of χ and f, where the mismatch between $f_{ground\ truth}$ and $f(\chi_{predicted})$ should be taken into account.

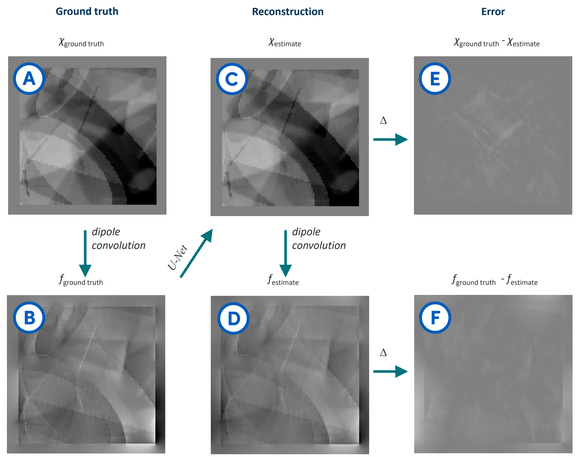

Figure 4. **Simulated evaluation data (left column), reconstruction results (middle column) and difference between ground truth and reconstruction (right column).** The retrieved susceptibility (C) is almost in indistinguishable from the ground truth (A).

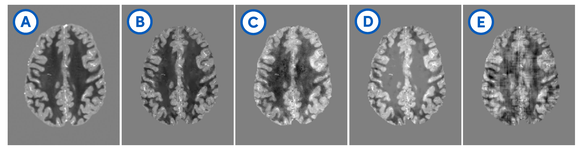

Figure 5. **Different reconstructions using an identical network architecture, but different training patterns.** A: Ground truth susceptibility distribution of one of the 2019 QSM Challenge datasets. B: Solution using the proposed training strategy. The white matter still shows false inhomogeneities. C: Solution when training without artificial noise showing noise amplification. D: Solution when the training patterns are too small, i.e. have a too large proportion of zero background (see Fig. 6) showing poor reconstruction of large white matter regions. E: Solution when no source-free area is simulated in the training data at the edge of the samples showing vertical and horizontal stripe artifact.

Acknowledgements

This work was supported by the German Federal Ministry of Education and Research (BMBF) within the TeleBrain project (01DS19009A), the Free State of Thuringia within the ThiMEDOP project (2018 IZN 0004) with funds of the European Union (EFRE), and the National Center for Advancing Translational Sciences of the National Institutes of Health under Award Number UL1TR001412. The content is solely the responsibility of the authors.

References

[1] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *arXiv:1505.04597 [cs]*, May 2015, Accessed: Aug. 17, 2018. [Online]. Available: http://arxiv.org/abs/1505.04597.
[2] S. Bollmann *et al.*, "DeepQSM - using deep learning to solve the dipole inversion for quantitative susceptibility mapping," *NeuroImage*, vol. 195, pp. 373–383, Jul. 2019, doi: 10.1016/j.neuroimage.2019.03.060.
[3] J. Yoon *et al.*, "Quantitative susceptibility mapping using deep neural network: QSMnet," *NeuroImage*, vol. 179, pp. 199–206, Oct. 2018, doi: 10.1016/j.neuroimage.2018.06.030.
[4] J. P. Marques and R. Bowtell, "Application of a Fourier-based method for rapid calculation of field inhomogeneity due to spatial variation of magnetic susceptibility," *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, vol. 25B, no. 1, pp. 65–78, Apr. 2005, doi: 10.1002/cmr.b.20034.
[5] F. Schweser and R. Zivadinov, "Quantitative susceptibility mapping (QSM) with an extended physical model for MRI frequency contrast in the brain: a proof-of-concept of quantitative susceptibility and residual (QUASAR) mapping," *NMR in Biomedicine*, Sep. 2018, doi: 10.1002/nbm.3999.
[6] T. Jochmann, J. Haueisen, R. Zivadinov, and F. Schweser, "U2-Net for DEEPOLE QUASAR–A Physics-Informed Deep Convolutional Neural Network that Disentangles MRI Phase Contrast Mechanisms," presented at the International Society of Magnetic Resonance in Medicine (ISMRM) 27th Annual Meeting, Montreal, Canada, May 13, 2019.
[7] H. Gudbjartsson and S. Patz, "The rician distribution of noisy mri data," *Magnetic Resonance in Medicine*, vol. 34, no. 6, pp. 910–914, 1995, doi: 10.1002/mrm.1910340618.

BNAC MBL
Buffalo Neuroimaging Analysis Center
MRI Biophysics Lab

mbl.bnac.net

Corresponding author: F.S. (schweser@buffalo.edu)