

Deep generative MRI reconstruction for unsupervised Gibbs ringing correction

Lucilio Cordero-Grande^{1,2}, Enrique Ortuño-Fisac¹, Jonathan O'Muircheartaigh^{2,3}, Jo Hajnal², and María Jesús Ledesma-Carbayo¹

¹Biomedical Image Technologies, ETSI Telecomunicación, Universidad Politécnica de Madrid & CIBER-BBN, Madrid, Spain, ²Centre for the Developing Brain & Biomedical Engineering Department, School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom, ³Department for Forensic and Neurodevelopmental Sciences, Institute of Psychiatry, Psychology and Neuroscience & MRC Centre for Neurodevelopmental Disorders, King's College London, London, United Kingdom

Synopsis

MRI reconstruction is formulated as the retrieval of the parameters of a deep decoder network fitted to the observations by an image formation model including the truncation of high-frequency information. Solutions without Gibbs ringing at any prescribed image grid can be obtained naturally by the model without training or ad-hoc post-corrections. We present quantitative and visual results for spectral extrapolation of magnitude images at different scales in an in-silico experiment and a high resolution ex-vivo brain MRI scan. After minor modifications to deal with complex data, the architecture is applied to 2D parallel imaging showing promising visual results.

Introduction

Truncation of high-frequency information above the maximum spatial frequency sampled during the Magnetic Resonance Imaging (MRI) experiment provokes spurious Gibbs Ringing (GR) next to tissue transitions. This may degrade the radiological quality of the images and confound segmentation algorithms, especially when using sinc interpolation. Classical solutions involve the application of a low pass filter to attenuate k-space discontinuities, but they introduce blurring. Recent works have explored spectral extrapolation¹ by deep learning, including self-supervised methods^{2,3}, knowledge transfer⁴, and combination with previous non-linear approaches⁵. However, these techniques are limited by availability of training data or transferability of results to different scales or sequences. The Deep Decoder⁶ (DD) has been introduced to solve inverse imaging problems without training, relaxing computational requirements of alternative approaches⁷, and offering competitive results with respect to training schemes in its application to MRI reconstruction⁸. Its functioning is based on implicitly biasing the reconstruction results towards the statistics of natural images, which could be particularly effective to circumvent GR.

Methods

Image inversion based on the DD generative model is achieved by:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \|\mathbf{M}\mathbf{A}\mathcal{G}_{\theta}(\mathbf{z}) - \mathbf{y}\|_2^2 \text{ with } \hat{\mathbf{x}} = \mathcal{G}_{\theta}(\mathbf{z})$$

with \mathbf{M} a mask indicating the acquired k-space samples, \mathbf{A} a forward model matrix, \mathcal{G} the DD generator, \mathbf{z} a fixed random input vector, θ the parameters of the network, to be fitted by backpropagation, \mathbf{y} the padded observations, and \mathbf{x} the reconstruction, obtained by propagating \mathbf{z} through \mathcal{G} using the fitted parameters $\hat{\theta}$. We study two models for \mathbf{A} :

$$\mathbf{A} = \mathbf{F},$$

with \mathbf{F} the Fourier transform of the image, used for GR correction from magnitude images, and

$$\mathbf{A} = \mathbf{F}\mathbf{O}\mathbf{S},$$

with \mathbf{O} a spatial overlap matrix from uniformly subsampled k-space, and \mathbf{S} a coil sensitivity matrix, used for direct spectral extrapolation of complex reconstructions from accelerated k-space data in parallel imaging. Results are provided for 2D reconstructions.

The original DD architecture has been modified for improved performance⁹ (Fig. 1). In our experiments, we use $L = \operatorname{ceil}(\log_2 \sqrt{N_1 N_2}) - 3$ DD scales with N_1 and N_2 the target grid sizes for \mathbf{x} , a $N_{\theta} \simeq N_1 N_2$ parameter network, and the Adam optimizer with 20000 iterations.

The modified Shepp-Logan phantom of size 512×512 is used as Ground Truth (GT) for in-silico validation of GR correction when truncating the original spectrum to 1/8, 1/4, and 1/2 the original sizes for both axes. A publicly available 0.1 mm resolution ex-vivo scan of a human brain¹⁰ is used to study the GR problem for truncation at 0.5 mm, 1 mm and 2 mm by comparing reconstructions of an axial slice with GT datasets downsampled to twice the corresponding truncation frequencies. Peak Signal to Noise Ratio (PSNR), Structural Similarity Index Measure¹¹ (SSIM) and a specific metric termed Unsampling Information Retrieval Ratio (UIRR):

$$\text{UIRR} = 1 - \frac{\|(\mathbf{I} - \mathbf{M})\mathbf{F}(\hat{\mathbf{x}} - \mathbf{x}^*)\|_2^2}{\|(\mathbf{I} - \mathbf{M})\mathbf{F}\mathbf{x}^*\|_2^2},$$

with \mathbf{I} the identity matrix and \mathbf{x}^* the GT, are computed for both experiments. In all cases we compare the DD generative fit, Zero-Filled (ZF) images, and DD for Filling (DDF) the unsampled areas of the spectrum only. Finally, a volumetric Cartesian TSE brain sequence acquired on a consented volunteer using a Philips 3 T Achieva scanner (1 mm isotropic resolution, $240 \times 188 \times 240$ mm IS-LR-AP field of view, 1.4×1.4 acceleration, elliptical shutter, repetition time $T_R = 2.5$ s, and echo time $T_E = 254.3$ ms) is used to test the complex integrated formulation. Namely, parallel imaging generative reconstruction onto a 0.5 mm pixel size image at the center of the IS readout direction is compared to standard ZF SENSitivity Encoded (SENSE) reconstruction without and with GR mitigation based on a Tukey filter (taper ratio 0.2, $ZF_{0.2}$).

Results

Fig. 2 (results for 1/4 truncation) shows noticeable reduction of GR oscillations when comparing DD or DDF to ZF. Moreover, visual results and Power Spectral Density (PSD) plots averaged along the two image axes suggest prominent retrieval of Unsampling High Frequency (UHF) components for this piecewise continuous phantom. Visual improvement in GR suppression and contrast when using the DD is observed in the ex-vivo experiment in Fig. 3 (results for 1 mm truncation), but in this case richer UHF information is only partially retrieved. As reported in Fig. 4, DD reconstructions provide improved PSNR and UIRR over ZF for all compared configurations, with improvements in SSIM at 0.25 mm GT resolution observed only for the DDF strategy. Moreover, UIRR approaches 1 for the phantom and is noticeably higher than 0 for brain data. Reconstructions in Fig. 5 show ringing effects for ZF-SENSE, which are reduced for $ZF_{0.2}$ -SENSE at the price of blurred structures. The DD reconstruction provides a better trade-off with minor ringing and sharp definition of fine detailed structures.

Discussion

We have presented a deep generator based method for unsupervised GR correction that can be used as a post-processing step or integrated into classical reconstruction formulations. Results show effective correction of GR artifacts with partial retrieval of UHF components in a variety of scales and truncation factors. Comparisons with supervised methods and tests on partial Fourier reconstruction or different acquisition encodings remain to be performed. For application to 3D data we plan to study decomposition into patches and convergence acceleration to alleviate memory and run time requirements. These refinements are also expected to improve the DD fit for large grid sizes.

Conclusion

Unsupervised GR correction can be achieved with generative deep architectures, either as a post-processing step or by integration into the reconstruction, with flexible applicability at different scales or to different sequences.

Acknowledgements

This work is funded by the Ministry of Science and Innovation, Spain, under the Beatriz Galindo Programme [BGP18/00178]. This work has been supported by the Madrid Government (Comunidad de Madrid-Spain) under the Multiannual Agreement with Universidad Politécnica de Madrid in the line Support for R&D projects for Beatriz Galindo researchers, in the context of the V PRICIT (Regional Programme of Research and Technological Innovation).

References

- Candès EJ, Fernandez-Granda C. Towards a mathematical theory of super-resolution. *Comm Pure Appl Math*, LXVII:0906-0956, 2014.
- Zhang Q, Ruan G, Yang W, Liu Y, Zhao K, Feng Q, Chen W, Wu EX, Feng Y. MRI Gibbs ringing artifact reduction by means of machine learning using convolutional neural networks. *Magn Reson Med*, 82:2133-2145, 2019.
- Muckley MJ, Ades-Aron B, Papaioannou A, Lemberskiy G, Solomon E, Lui YW, Sodickson DK, Fieremans E, Novikov DS, Knoll F. Training a neural network for Gibbs and noise removal in diffusion MRI. *Magn Reson Med*, 85:413-428, 2021.
- Zhao X, Zhang H, Zhou Y, Bian W, Zhang T, Zou X. Gibbs-ringing artifact suppression with knowledge transfer from natural images to MR images. *Multimed Tools Appl*, 79: 33711-33733, 2020.
- Penkin MA, Krylov AS, Khvostikov AV. Hybrid method for Gibbs-ringing artifact suppression in magnetic resonance images. *Program Comput Softw*, 47(3):207-214, 2021.
- Heckel R, Hand P. Deep decoder: concise image representations from untrained non-convolutional networks. *ICLR*, 2019.
- Ulyanov D, Vedaldi A, Lempitsky V. Deep image prior. *Int J Comput Vis*, 128:1867-1888, 2020.
- Darestani MZ, Heckel R. Accelerated MRI with un-trained neural networks. *IEEE Trans Comput Imaging*, 7:724-733, 2021.
- Cordero-Grande L, Ortuño-Fisac JE, Uus A, Deprez M, Santos A, Hajnal JV, Ledesma-Carbayo MJ. Fetal MRI by robust deep generative prior reconstruction and diffeomorphic registration: application to gestational age prediction. *arXiv:2111.00102*, 2021.
- Edlow BL, Mareyam A, Horn A, Polimeni JR, Witzel T, Tisdall MD, Augustinack JC, Stockmann JP, Diamond BR, Stevens A, Tirrell LS, Folkerth RD, Wald LL, Fischl B, van der Kouwe A. 7 Tesla MRI of the ex vivo human brain at 100 micron resolution. *Sci Data*, 6:244, 2019.
- Wang Z, Bovik, AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process*, 13(4):600-612, 2004.

Figures




Fig. 1. DD architecture used as image generator. The example shows a network with $L=5$ scales for fitting a $C=1$ -channel magnitude image of size $N=[N_1, N_2]$. We use a scaling ratio of 2 and same number of channels K for all scales. The coarsest scale comprises simply a linear layer. The remaining scales are connected by sinc upsampling layers and formed by blocks of batch normalization, activation and linear layers. At the finest scale, we use two such blocks followed by sigmoid (magnitude, $C=1$) or hyperbolic tangent (complex, $C=2$) activation.




Fig. 2. Visual deringing results for the modified 512x512 Shepp-Logan phantom at truncation factors 4x4 (see sampling mask M for a 512x512 k-space region). Top row: results. Bottom row: differences with the GT. We observe decreased errors for DD or DDF when compared to ZF, with suppression of ringing and prominent retrieval of UHR features, which is supported by the PSD curves (DD and DDF almost overlapping with GT for the unsampled k-space frequencies).




Fig. 3. Visual deringing results for the ex-vivo dataset for reconstructing to a 0.5mm grid with truncation at 1mm (see sampling mask M for a 0.5mm k-space region). Top row: results. Bottom row: differences with the GT. We observe decreased errors for DD or DDF when compared to ZF, with suppression of ringing and partial retrieval of UHR features as shown in the zoomed areas (see also the PSD curves).

Truncation	Phantom				Ex-vivo				
	Metric	ZF	DD	DDF	Resolution	Metric	ZF	DD	DDF
2×2	PSNR	28.80	53.74	54.23	0.5 \rightarrow 0.25	PSNR	33.53	34.28	34.65
	SSIM	0.930	0.999	0.999		SSIM	0.880	0.967	0.889
	UIRR	0.000	0.007	0.000		UIRR	0.000	0.227	0.000
	TUHR	0.000	0.000	0.000		TUHR	0.000	0.227	0.000
4×4	PSNR	25.54	42.93	42.96	1.0 \rightarrow 0.5	PSNR	33.13	34.04	34.99
	SSIM	0.960	0.996	0.996		SSIM	0.900	0.912	0.945
	UIRR	0.000	0.002	0.000		UIRR	0.000	0.348	0.000
	TUHR	0.000	0.002	0.000		TUHR	0.000	0.348	0.000
8×8	PSNR	21.78	34.28	34.28	2.0 \rightarrow 1.0	PSNR	31.32	32.19	32.50
	SSIM	0.810	0.986	0.986		SSIM	0.829	0.844	0.844
	UIRR	0.000	0.944	0.000		UIRR	0.000	0.248	0.000
	TUHR	0.000	0.944	0.000		TUHR	0.000	0.248	0.000

Fig. 4. Quantitative deringing results for the modified 512x512 Shepp-Logan phantom at different truncation factors and for the ex-vivo dataset with 2x2 truncation at different scales (resolution in mm). PSNR (in dBs) and UIRR (reference value of 0 for ZF) are improved in all experiments when using the DD or DDF as compared to ZF. SSIM does generally improve when using the DD or DDF with the exception of the finest scale in the ex-vivo experiment for the DD case. This is attributed to incomplete convergence of the DD fit.




Fig. 5. Results for the in-vivo dataset for reconstructing to a 0.5mm grid with elliptical truncation at 1mm (see sampling mask M for a 0.5mm k-space region). The DD result provides a good trade-off between suppression of ringing present in ZF and avoidance of blurring in $ZF_{0.2}$. Traces of UHF retrieval may be observed in specific interfaces (see the extra-axial cerebrospinal fluid in the zoomed areas).