GSURE Denoising enables training of higher quality generative priors for accelerated Multi-Coil MRI Reconstruction

Asad Aali[1], Marius Arvinte[1,2], Sidharth Kumar[1], Yamin Ishraq Arefeen[1], and Jonathan I. Tamir[1]
[1]Chandra Family Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX, United States, [2]Intel Corporation, Hillsboro, OR, United States

Synopsis

Keywords: AI/ML Image Reconstruction, Image Reconstruction, Deep Generative Models, Inverse Problems, Unsupervised Learning, Denoising

Motivation: Publicly available k-space data used for training are inherently noisy with no available ground truth.

Goal(s): To denoise k-space data in an unsupervised manner for downstream applications.

Approach: We use Generalized Stein's Unbiased Risk Estimate (GSURE) applied to multi-coil MRI to denoise images without access to ground truth. Subsequently, we train a generative model to show improved accelerated MRI reconstruction.

Results: We demonstrate: (1) GSURE can successfully remove noise from k-space; (2) generative priors learned on GSURE-denoised samples produce realistic synthetic samples; and (3) reconstruction performance on subsampled MRI improves using priors trained on denoised images in comparison to training on noisy samples.

Impact: This abstract shows that we can denoise multi-coil data without ground truth and train deep generative models directly on noisy k-space in an unsupervised manner, for improved accelerated reconstruction.

Introduction

Recent work has leveraged large-scale medical imaging datasets to train generative priors used to solve ill-posed inverse problems. A common assumption is that the data available are noise-free and fully sampled, but this can be prohibitively expensive to collect in most cases. Even large-scale fully sampled datasets, such as FastMRI [1] are inherently noisy, yet they are used to train generative priors without accounting for this noise[2–4]. Stein's unbiased risk estimate (SURE) represents an unbiased training method that learns to denoise Gaussian noise-corrupted samples in a self-supervised manner, without access to clean ground truth [5,6]. Generalized SURE (GSURE) represents an extension of SURE[7] for training a model using samples corrupted by noise from the exponential family which includes the application of any linear operator. Here we apply GSURE to fully sampled, noisy multi-coil k-space to remove the noise inherent in the data without access to a ground-truth image. We then use this denoised dataset to train a score-based deep generative prior that is used for subsampled MRI reconstruction.

Methods

Multicoil MRI data are acquired in k-space, with measurements denoted by $y = FSx + n$, where F is the Fourier transform, S is the coil sensitivity operator[8], x is the ground-truth, complex-valued image, and n is the sensor noise in k-space and is assumed to be Gaussian with non-white covariance matrix. In low-field imaging applications with low SNR, this noise can be substantially limiting[9]. Our premise is that by taking this noise into account we can formulate and learn an improved generative prior from denoised measurements, and subsequently use this prior in subsampled MRI reconstruction.

We arrive at $y^* = FSx^* + w$, where $w \sim N(0, I)$:
We first apply noise pre-whitening to the fully sampled multi-coil k-space using BART[10,11]. We (b) apply the GSURE loss function [7] to train a denoiser without ground-truth, that can be used to estimate a coil-combined denoised image. The GSURE denoiser takes as input the adjoint of the measurements scaled by the estimated noise variance. The GSURE loss matches the supervised denoising loss function in expectation, as proved by[7]. g_{phi} is the learned denoising model, and we approximate the divergence term in the loss function using monte-carlo SURE[12]. Using the denoised images, (c) we train a score-based generative model using denoising score-matching as outlined in [13]. We then reconstruct subsampled MRI data through posterior sampling via annealed Langevin dynamic.

To illustrate the effect, we use T2-weighted brain (average SNR 32 dB) and fat-suppressed proton-density knee (average SNR 24 dB) FastMRI data[1], which is inherently noisy. We apply our approach directly to this data, and after adding additional noise to show the impact at lower inherent SNR – 22 dB and 14 dB for the brain and knee, respectively. Using the brain MRI, we train "naive" score-based generative models on the inherently noisy data, and "GSURE-Score" models on the GSURE-denoised data. We retrospectively subsample 57 test samples with an acceleration factor of 5 and compare reconstruction quality using the various deep generative priors. We use normalized root mean squared error (NRMSE) for quantitative evaluation.

Results and Discussion

Figs 2 and 3 show the denoising results with networks trained in a supervised manner and with the proposed GSURE denoising. It can be observed that GSURE denoising preserves the signal while substantially reducing the noise power, in line with the supervised denoisers. At lower SNRs, the effect of GSURE-denoising is more pronounced, which is an important consideration for low-SNR imaging such as in low-field MRI.

Fig 4 shows the prior samples generated from score-based generative models naively trained on noisy data versus GSURE-denoised data. At high-training SNR, the quality of prior samples remains consistent. At lower training SNR, the naive model generates unrealistic images.

Fig 5 shows the result of reconstruction from R=5 retrospectively subsampled data. The reconstruction using the score-based generative model trained on GSURE-denoised data shows quantitative and qualitative improvement over the naively trained models, where the effect is most pronounced when the training set has a lower SNR. This holds true over the 57 test samples.

Conclusions and Future Directions

We showed the benefit of noise pre-whitening and unsupervised denoising of publicly available k-space before training deep learning-based reconstruction methods. In addition to denoising multi-coil data without ground truth, our approach provides a mechanism for training more faithful score-based generative priors that can be used for downstream applications such as image reconstruction. Substantial improvement over naively trained score-based models promotes further development of models that can deal with noise corruption of different types. This can also be extended towards learning directly from undersampled measurements in future works[14,15] .

Acknowledgements

This work was supported by NSF IFML 2019844 and NSF CCF-2239687 (CAREER).

References

1. J. Zbontar et al., “fastMRI: An open dataset and benchmarks for accelerated MRI,” 2018, arXiv:1811.08839. [Online]. Available: http://arxiv.org/abs/1811.08839
2. Song, Y., Shen, L., Xing, L., & Ermon, S. (2021). Solving inverse problems in medical imaging with score-based generative models. ICLR 2022.
3. H. Chung, E. S. Lee, and J. C. Ye, “MR image denoising and superresolution using regularized reverse diffusion,” IEEE Transactions on Medical Imaging,vol. 42, no. 4, pp. 922–934, Apr. 2023.
4. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A. G., & Tamir, J. (2021). Robust compressed sensing mri with deep generative priors. Advances in Neural Information Processing Systems, 34, 14938-14954.
5. Asad Aali, Marius Arvinte, Sidharth Kumar, and Jonathan I Tamir. Solving inverse problems with score-based generative priors learned from noisy data. in Proceedings of Asilomar Conference on Signals, Systems, and Computers, 2023.
6. C. M. Stein, “Estimation of the mean of a multivariate normal distribution,” The annals of Statistics, pp. 1135–1151, 1981.
7. Yonina C Eldar. Generalized SURE for exponential families: Applications to regularization. IEEE Transactions on Signal Processing, 57(2):471–481, 2008.
8. Uecker Martin, Lai Peng, Murphy Mark J., Virtue Patrick, Elad Michael, Pauly John M., Vasanawala Shreyas S., Lustig Michael. ESPIRiT-an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA. Magnetic Resonance in Medicine. 2014;71 (3): 990–1001. http://dx.doi.org/10.1002/mrm.24751.

9. Marques JP, Simonis FFJ, Webb AG. Low-field MRI: An MR physics perspective. J Magn Reson Imaging. 2019 Jun;49(6):1528-1542. doi: 10.1002/jmri.26637. Epub 2019 Jan 13. PMID: 30637943; PMCID: PMC6590434.
10. Martin Uecker, Frank Ong, Jonathan I Tamir, Dara Bahri, Patrick Virtue, Joseph Y Cheng, Tao Zhang, and Michael Lustig. Berkeley advanced reconstruction toolbox. In Proc. Intl. Soc. Mag. Reson. Med, volume 23, 2015.
11. Kellman P, McVeigh ER. Image reconstruction in SNR units: a general method for SNR measurement. Magn Reson Med. 2005 Dec;54(6):1439-47. doi: 10.1002/mrm.20713. Erratum in: Magn Reson Med. 2007 Jul;58(1):211-2. PMID: 16261576; PMCID: PMC2570032.
12. S. Ramani, T. Blu, and M. Unser, "Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms," IEEE Transactions on image processing, vol. 17, no. 9, pp. 1540–1554, 2008.
13. Song, Y., & Ermon, S. (2020). Improved techniques for training score-based generative models. Advances in neural information processing systems, 33, 12438-12448.
14. Aggarwal, H. K., Pramanik, A., John, M., & Jacob, M. (2022). ENSURE: A general approach for unsupervised training of deep image reconstruction algorithms. IEEE Transactions on Medical Imaging, 42(4), 1133-1144.
15. Kawar, B., Elata, N., Michaeli, T., & Elad, M. (2023). GSURE-Based Diffusion Model Training with Corrupted Data. arXiv preprint arXiv:2305.13128.

Figures

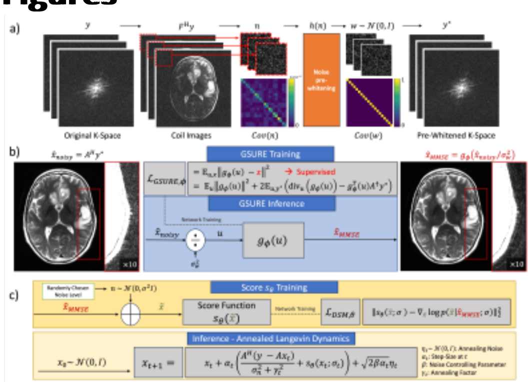

Fig 1 (a) Illustration of FastMRI noise estimation and pre-whitening from fully sampled multi-coil k-space samples. (b) After pre-whitening, a denoising network is trained in an unsupervised manner directly from the noisy data using GSURE, which is equivalent in expectation to supervised learning. (c) Using the denoised data, a score-based generative model is trained using denoising score matching.

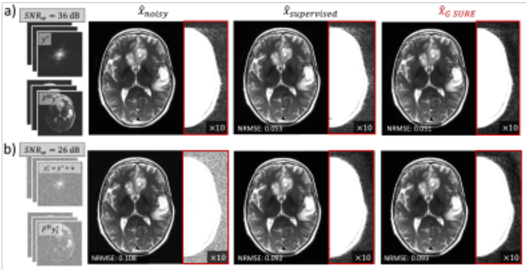

Fig. 2: Brain FastMRI denoising results comparing supervised and GSURE-based training at two SNR values. The first column for each SNR value depicts the adjoint image of pre-whitened k-space data normalized by the 99th percentile of the RSS image. The second column shows the denoised image obtained through a supervised denoiser network as shown in Fig.1. Third column shows the GSURE denoiser network trained in a self-supervised manner without access to the ground truth image. Panel (b) further adds noise to investigate the low-SNR setting.

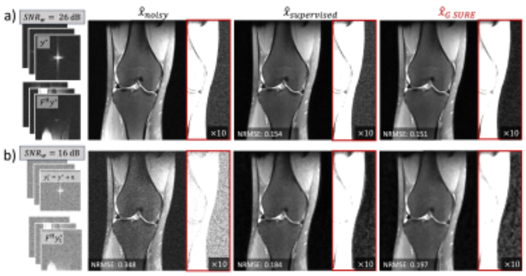

Fig. 3: Knee FastMRI denoising results comparing supervised and GSURE-based training at two SNR values. The first column for each SNR value depicts the adjoint image of pre-whitened k-space data normalized by the 99th percentile of the RSS image. The second column shows the denoised image obtained through a supervised denoiser network as shown in Fig.1. Third column shows the GSURE denoiser network trained in a self-supervised manner without access to the ground truth image. Panel (b) further adds noise to investigate the low-SNR setting.

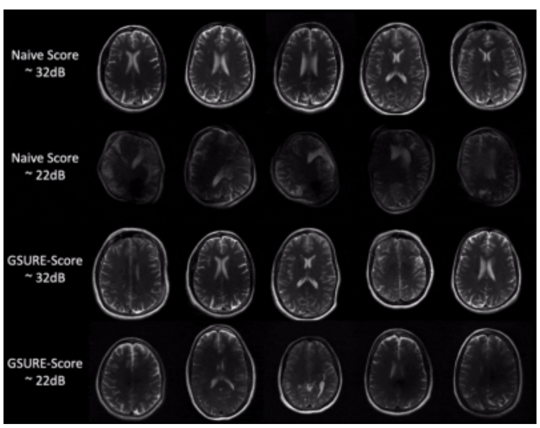

Fig. 4 Sample MR images generated from the deep generative priors trained at different SNR levels and with and without GSURE-denoising as a preprocessing step.

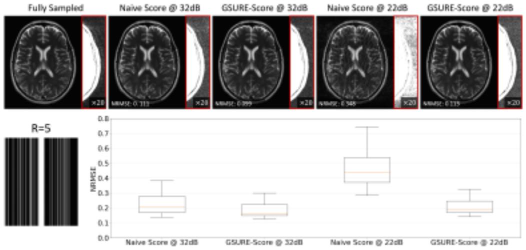

Fig. 5 Retrospective reconstruction results at R=5 using score-based generative models show improved reconstruction performance with GSURE-

denoising, even when the data are at high SNR. NRMSE values over 57 validation samples with naive and GSURE-Score for posterior sampling.