2021 Online International Conference on Computational Harmonic Analysis (Online-ICCHA2021)

Provably Accurate, Stable and Efficient Deep Neural Networks for Compressive Imaging

Maksym Neyra-Nesterenko
Department of Mathematics
mneyrane@sfu.ca

Ben Adcock
Department of Mathematics
ben_adcock@sfu.ca

***Abstract*—Machine learning, and in particular, Deep Learning (DL), has recently demonstrated substantial potential to outperform standard methods in the field of compressive imaging; that is, the reconstruction of images from highly undersampled measurements. Yet, unlike standard methods, DL approaches generally lack theoretical guarantees on the accuracy and stability of reconstructing images. Moreover, it has also been well documented that certain DL approaches can be unstable and/or exhibit poor, inconsistent or otherwise unexpected generalization. This raises the question: under what conditions can provably accurate, stable and efficient Deep Neural Networks (DNNs) be computed for compressive imaging problems? In this work, inspired by recent work, we answer this question in the affirmative in the case of subsampled Fourier measurements. Specifically, we show the existence of efficient DNNs that are stable and accurate for gradient-sparse images, thus complementing and extending a recent result for Haar wavelet-sparse images. This work helps confirm the significant potential of DL, and may shed light onto how its practical performance can be further improved.**

I. Compressive imaging

Reconstructing images from highly undersampled measurements is a key task in many medical, scientific and industrial applications, with important examples being Magnetic Resonance Imaging (MRI), X-Ray Computed Tomography (CT), electron microscopy and seismic imaging, to name but a few. Typically an image reconstruction problem is formulated as the discrete linear inverse problem: given $y = Ax + e \in \mathbb{C}^m$, recover $x \in \mathbb{C}^N$. Here $x \in \mathbb{C}^N$ is the unknown image, $y \in \mathbb{C}^m$ are the noisy measurements and $A \in \mathbb{C}^{m \times N}$ is the measurement matrix. This problem is ill-posed, since $m \ll N$ in practice. Yet over the last two decades, nonlinear reconstructions techniques based on sparse regularization have given rise to significant performance improvements over classical linear reconstruction methods. Among the most popular is TV-minimization (see, e.g., [1], [2]), wherein a reconstruction $\hat{x}$ is computed by solving

$$\min_{z \in \mathbb{C}^N} \|z\|_{\mathrm{TV}} \quad \text{subject to } \|Az - y\|_{\ell^2} \leq \eta, \tag{1}$$

where $\|z\|_{\mathrm{TV}} = \|\nabla z\|_{\ell^1}$ is the TV-norm. Along with wavelet-based methods, TV minimization is a standard benchmark procedure for compressive imaging. Further, both approaches have well-understood theoretical performance (stability and accuracy) guarantees under suitable conditions on A [2].

II. Deep learning for compressive imaging

In the last five years, new techniques for compressive imaging based on DL have emerged [3]. Here, one uses training data, i.e. pairs $\{(y_i, x_i)\}_{i=1}^k$, where $y_i = Ax_i + e_i$, to learn a reconstruction map $\mathcal{N} : \mathbb{C}^m \to \mathbb{C}^N$, which typically takes the form of a DNN. Various different approaches have been introduced to do this, such as learn-the-physics methods, denoiser methods, unravelled optimization solvers, learned proximal maps, plug-and-play methods, and various others. See, for example, [4], [2], [5], [6], [7], [8], [9] for overviews.

While these methods have exhibited superior performance over standard techniques, there is increasing awareness that they can suffer from key drawbacks. They have a tendency to be unstable to small perturbations [10], [11], [12], [13], and to exhibit inconsistent generalization behaviours, referred to as AI-generated *hallucinations* [13]. Notably, such behaviours are typically absent in standard methods. Theoretical justification for why DL can lead to this behaviour has been presented in [14].

It is therefore critical to catergorize, understand and, ultimately, improve the robustness of DNN-based methods. One key strand of this endeavour involves studying the question of when accurate, stable and efficient DNNs can be computed. Recently, this question was studied in [15], yielding both negative (uncomputability) results and positive results, the latter showing that DNNs can match the stability and accuracy of Haar wavelet-based sparse regularization. In this work, we complement and extend this study by examining the case of gradient-based sparse recovery methods such as TV-minimization.

III. Main result

We consider Fourier imaging, in which $A = m^{-1/2} P_\Omega F$ is a subsampled Fourier matrix according to some subsampling pattern $\Omega \subseteq \{1, \dots, N\}$, $|\Omega| = m$. Fix $\eta \geq 0$, $2 \leq s \leq N$, $\chi > 0$ and consider the class $\mathbb{I} = \mathbb{I}_\chi$ given by

$$\mathbb{I} = \left\{ (x, e) \in \mathbb{C}^N \times \mathbb{C}^m : \|x\|_{\ell^2} \leq 1,\ \|e\|_{\ell^2} \leq \eta,\ \mathcal{CS}(x, \eta) \leq \chi \right\},$$

where $\mathcal{CS}(x, \eta) := \frac{\sigma_s(\nabla x)_{\ell^1}}{\sqrt{s}} + \sqrt{\log N}\eta$. Now let

$$\mathbb{M} = \mathbb{M}_{\Omega,\chi} = \{ y = m^{-1/2} P_\Omega F x + e : (x, e) \in \mathbb{I} \}.$$

Theorem 3.1 (Stable, accurate and efficient DNNs): There is a distribution on Ω for which the following holds, provided $m \gtrsim s\,\mathrm{polylog}(s, N, \epsilon)$ for any $0 < \epsilon < 1$. There exists a neural network $\mathcal{N} : \mathbb{C}^m \to \mathbb{C}^N$ whose maximum width is at most a constant times $N^2 + m$ and whose depth is proportional to $\log(\chi^{-1})$, for which, with probability at least $1 - \epsilon$,

$$\|x - \mathcal{N}(y)\|_{\ell^2} \lesssim \chi, \quad \forall y = Ax + e \in \mathbb{M}_{\Omega,\chi}.$$

Observe that $\mathbb{M}$ is the class of noisy measurements of images that can be recovered up to error χ via (1), i.e. stably (due to the η term) and accurately (due to $\sigma_s(\nabla x)_{\ell^1}$) [16], [17], [18]. Hence, this result states that a DNN can achieve the same performance; in particular, stability and predictable generalization near the set of gradient-sparse images. Further, this network is computationally efficient, since its depth scales only logarithmically with χ.

The proof of this theorem is completely constructive. It employs a novel architecture based on an unravelled NESTA solver [19] in combination with a restart scheme [20] to provide exponential convergence, plus a novel sampling strategy based on Bernoulli selectors and compressed sensing analysis. It highlights the vast potential of DL, and may shed light on how to train better DNNs without succumbing to effects such as instabilities and AI-generated hallucinations.

References

[1] A. Chambolle, M. Novaga, D. Cremers, and T. Pock, "An introduction to total variation for image analysis," in *Theoretical Foundations and Numerical Methods for Sparse Recovery*, ser. Radon Series in Computational and Applied Mathematics, M. Fornasier, Ed. Berlin: de Gruyter, 2010, vol. 9, pp. 263–340.

[2] B. Adcock and A. C. Hansen, *Compressive Imaging: Structure, Sampling, Learning*. Cambridge: Cambridge University Press (in press), 2021.

[3] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, "Image reconstruction is a new frontier of machine learning," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1289–1296, 2018.

[4] M. T. McCann, K. H. Jin, and M. Unser, "Convolutional neural networks for inverse problems in imaging: a review," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 85–95, 2017.

[5] G. Ongie, A. Jalal, C. A. Metzler, R. G. Baraniuk, A. G. Dimakis, and R. Willett, "Deep learning techniques for inverse problems in imaging," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 39–56, 2020.

[6] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, "Solving inverse problems using data-driven models," *Acta Numer.*, vol. 28, pp. 1–174, 2019.

[7] S. Ravishankar, J. C. Ye, and J. A. Fessler, "Image reconstruction: from sparsity to data-adaptive methods and machine learning," *Proc. IEEE*, vol. 108, no. 1, pp. 86–109, 2020.

[8] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: interpretable, efficient deep learning for signal and image processing," *arXiv:1912.10557*, 2019.

[9] H. Zhang and B. Dong, "A review on deep learning in medical image reconstruction," *J. Oper. Res. Soc. China (in press)*, 2020.

[10] V. Antun, F. Renna, C. Poon, B. Adcock, and A. C. Hansen, "On instabilities of deep learning in image reconstruction and the potential costs of AI," *Proc. Natl. Acad. Sci. USA*, vol. 117, no. 48, pp. 30 088–30 095, 2020.

[11] Y. Huang, T. Würfl, K. Breininger, L. Liu, G. Lauritsch, and A. Maier, "Some investigations on robustness of deep learning in limited angle tomography," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2018, pp. 145–153.

[12] F. Knoll, T. Murrell, A. Sriram, N. Yakubova, J. Zbontar, M. Rabbat, A. Defazio, M. J. Muckley, D. K. Sodickson, C. L. Zitnick, *et al.*, "Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge," *Magn. Reson. Med.*, vol. 84, no. 6, pp. 3054–3070, 2020.

[13] M. J. Muckley, B. Riemenschneider, A. Radmanesh, S. Kim, G. Jeong, J. Ko, Y. Jun, H. Shin, D. Hwang, M. Mostapha, *et al.*, "State-of-the-art machine learning MRI reconstruction in 2020: results of the second fastMRI challenge," *arXiv:2012.06318*, 2020.

[14] N. M. Gottschling, V. Antun, B. Adcock, and A. C. Hansen, "The troublesome kernel: why deep learning for inverse problems is typically unstable," *arXiv:2001.01258*, 2020.

[15] M. J. Colbrook, V. Antun, and A. C. Hansen, "Can stable and accurate neural networks be computed? – on the barriers of deep learning and smale's 18th problem," *arXiv:2101.08286*, 2021.

[16] F. Krahmer and R. Ward, "Stable and robust sampling strategies for compressive imaging," *IEEE Trans. Image Process.*, vol. 23, no. 2, pp. 612–622, 2013.

[17] C. Poon, "On the role of total variation in compressed sensing," *SIAM J. Imaging Sci.*, vol. 8, no. 1, pp. 682–720, 2015.

[18] B. Adcock, N. Dexter, and Q. Xu, "Improved recovery guarantees and sampling strategies for tv minimization in compressive imaging," *SIAM J. Imaging Sci. (to appear)*, 2021.

[19] S. Becker, J. Bobin, and E. J. Candès, "NESTA: A fast and accurate first-order method for sparse recovery," *SIAM J. Imaging Sci.*, vol. 4, no. 1, pp. 1–39, 2011.

[20] V. Roulet and A. Boumal, N. d'Aspremont, "Computational complexity versus statistical performance on sparse recovery problems," *Inf. Inference*, vol. 9, no. 1, pp. 1–32, 2020.