We thank all reviewers for their comments. We apologize for the typos; all these minor points will be dealt with in the revised version of our article.

R#2: My only relatively minor issue with the paper...range of the earlier layers)..

Indeed, we did not clearly explain what exactly the lower bound proves, in two different ways:

- As you point out, we incorrectly asserted that it suffices to consider a one-layer neural network. Our lower bound does only apply to one-layer neural networks. If there are many layers, some expansive and some not, then things might still be fine; we do not know and it is indeed an interesting open problem.
- Our lower bound is only against recovering the latent parameter x^*, not against recovering the image $G(x^*)$. For recovering the image, it's not clear (to us, at least) that any expansion should be necessary – it's not information-theoretically necessary, and although global landscape analysis probably fails without expansion, there may be other algorithms.

Thank you for this interesting comment, we will clarify/fix these issues in the final version of our paper.

I think that in Line 84, the authors should add a reference to "A Geometric Analysis of Phase Retrieval". We believe that we have a reference to this exact paper on line 84 :-).

R#3: We really appreciate your positive feedback!

R#4: The signal model considered in the paper is essentially of theoretical interest: the signal is the output of a neural network with random weights, encoded by a low-dimensional vector at the input. Its connection to signals or measurement models of practical interest is unclear, which will limit the impact of the paper and its potential audience to very mathematically-oriented readers. The mathematical contribution may be interesting in itself but its applications seem somewhat restricted in scope.

The generative prior model (i.e. "signal is output of neural network with low latent dimension") is in fact of significant practical interest. It has been intensively studied in recent years in the context of compressed sensing, inpainting, and other image recovery problems [2]. Indeed, much empirical evidence suggests that the generative prior can enable image recovery with far fewer samples than sparsity-based priors (e.g. in the wavelet basis).

Our goal in this paper, and many papers before us [4, 3, 5], is to theoretically analyze the already established practical success of the generative prior model. Our contribution is to strengthen the existing results by making weaker assumptions. This brings us closer to rigorously explaining the practical success of this method.

Why Gaussian random weights? First, it's a common theoretical assumption, and success at analyzing Gaussians motivates study of more realistic distributions. Second, some works have shown that even trained neural networks have weights which look Gaussian in important ways (e.g. singular values of the weight matrices) [1]. Third, random initialization is a technique commonly used in practice; understanding the theoretical properties of a network at initialization is necessary to understand the properties after training. Indeed many theoretical works on overparametrized networks argue that the weights don't move a lot after training. Finally, when a trained GAN is unavailable, a practical approach, called Deep Image Prior is to take a network with randomly assigned weights, and use that as a regularizer. This is our setting. For further motivation for Gaussian weights, we refer to prior work in this area.

References

[1] Sanjeev Arora, Yingyu Liang, and Tengyu Ma. Why are deep nets reversible: A simple theory, with implications for training. *arXiv preprint arXiv:1511.05653*, 2015.

[2] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed sensing using generative models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 537–546. JMLR. org, 2017.

[3] Paul Hand, Oscar Leong, and Vlad Voroninski. Phase retrieval under a generative prior. In *Advances in Neural Information Processing Systems*, pages 9136–9146, 2018.

[4] Paul Hand and Vladislav Voroninski. Global guarantees for enforcing deep generative priors by empirical risk. *IEEE Transactions on Information Theory*, 66(1):401–418, 2019.

[5] Shuang Qiu, Xiaohan Wei, and Zhuoran Yang. Robust one-bit recovery via relu generative networks: Improved statistical rates and global landscape analysis. *arXiv preprint arXiv:1908.05368*, 2019.