Solving Inverse Imaging Problems with Generative Machine Learning Models

Margaret Duff
M.A.G.Duff@bath.ac.uk

Mathhias Ehrhardt
M.Ehrhardt@bath.ac.uk

Neill D. F. Campbell
N.Campbell@bath.ac.uk

Solving an inverse problem is the task of computing an unknown physical quantity from indirect measurements found via a forward model. Let $A : X \to Y$ be a forward process that takes an image $x \in X$ to data $y \in Y$. The inverse problem takes data, y corrupted by noise, ϵ, and finds image, x, such that $y = A(x) + \epsilon$. Interesting inverse problems are nearly always ill-posed. In particular, sometimes the data recovered is not sufficient to solve the reconstruction problem accurately and some form of prior knowledge is required.

Generative models learn, from observations, approximations to high-dimensional data distributions. The learnt distribution can then be used as a prior when solving the inverse problem, providing additional information and ensuring solutions are feasible.

A trained generative model $G : Z \to X$, takes values in a known lower dimensional latent space, Z and outputs feasible images. It is now possible to search over the latent space to solve the inverse problem:

$$z^* \in \arg\min_{z \in Z} ||A(G(z)) - y||, \quad x^* = G(z^*). \tag{1}$$

This method was first introduced by Bora et al. [1] for compressed sensing and has since been further explored [2]–[4]. In this work, we numerically evaluate and compare different generative models in a variety of applications. The generative models under investigation are primarily variational autoencoders (VAEs) [5] and generative adversarial networks (GANs) [6]. An inpainting example using a Wasserstein GAN is given in figure 1.

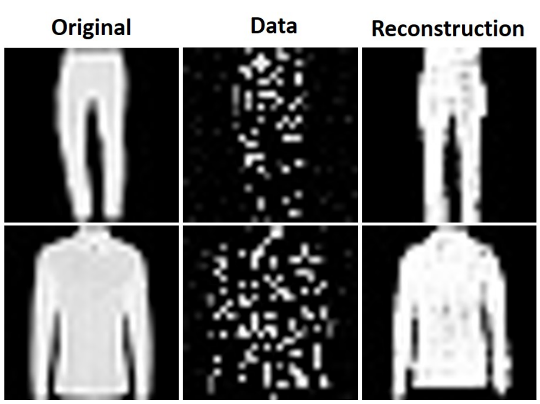

Fig. 1. A demonstration of using a Wasserstein GAN trained on the FMNIST [7] dataset for an inpainting problem. It shows reconstruction from noisy samples of just 25% of pixels.

Learnt priors could provide more specific information than a hand-crafted regulariser and thus a better reconstruction. Any learning would also not require paired training data and would be done independently of the forward model. This makes the method very flexible in real-world scenarios where noise levels and forward model parameters may change.

This approach raises a wide range of mathematical questions. The addition of a generator in (1) makes the minimisation problem potentially non-linear and non-convex. The need to add explicit regularisation to (1), to aid optimisation or to provide additional knowledge of the latent space, is currently unknown. Investigations into the nature of the learnt latent space and the mapping into the image space will provide insight into this problem, guiding the choice of regularisation and optimisation method. For example, a disconnected latent space will require many random restarts of an optimisation method and the generation of unfeasible images from sections of the latent space may suggest the need for additional regularisation. In contrast, smooth interpolations between images by moving through the latent space will aid the success of optimisation methods.

In a deterministic setting, we also require that the ground truth image lies in the range of the generator. Arora et al. [8] suggests that even well-known GAN approaches learn distributions with much lower support than the target distribution. This form of mode collapse is difficult to detect but is very important in applications where decision making depends on the reconstructed image.

Providing a set of desired criteria for evaluating a trained generative network, in the context of inverse problems, will allow comparison between generative model approaches and identify areas for future work.

References

[1] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed Sensing Using Generative Models," in *International Conference on Machine Learning*, 2017.
[2] V. Shah and C. Hegde, "Solving Linear Inverse Problems Using GAN Priors: An Algorithm with Provable Guarantees," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018.
[3] M. Dhar, A. Grover, and S. Ermon, "Modeling Sparse Deviations for Compressed Sensing Using Generative Models," in *International Conference on Machine Learning*, 2018.
[4] D. Obmann, J. Schwab, and M. Haltmeier, "Deep synthesis regularization of inverse problems," *arXiv preprint arXiv:2002.00155*, 2020.
[5] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *International Conference on Learning Representations*, 2014.
[6] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets,"
[7] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
[8] S. Arora, A. Risteski, and Y. Zhang, "Do GANs Learn the Distribution? Some Theory and Empirics," in *International Conference on Learning Representations*, 2018.