GAN Priors for Bayesian Inference

Dhruv V Patel
University of Southern California
dhruvvpa@usc.edu

Assad A Oberai
University of Southern California
aoberai@usc.edu

Abstract

Bayesian inference is used extensively to infer and to quantify the uncertainty in a field of interest from a measurement of a related field when the two are linked by a mathematical model. Despite its many applications, Bayesian inference faces challenges when inferring fields that have discrete representations of large dimension and/or have prior distributions that are difficult to characterize mathematically. In this work we demonstrate how the approximate distribution learned by a generative adversarial network (GAN) may be used as a prior in a Bayesian update to address both these challenges. We demonstrate the efficacy of this approach by inferring and quantifying uncertainty in a physics-based inverse problem and an inverse problem arising in computer vision. In this latter example, we also demonstrate how the knowledge of the spatial variation of uncertainty may be used to select an optimal strategy of placing the sensors (i.e. taking measurements), where information about the image is revealed one sub-region at a time.

1 Introduction

Bayesian inference is a principled approach to quntify uncertainty in inverse problems that are constrained by mathematical model (Kaipio and Somersalo [2006], Dashti and Stuart [2016], Polpo et al. [2018]). It has found applications in diverse fields such as geophysics (Gouveia and Scales [1997], Martin et al. [2012], Isaac et al. [2015]), climate modeling (Jackson et al. [2004]), chemical kinetics (Najm et al. [2009]), heat conduction (Wang and Zabaras [2004]), astrophysics (Loredo [1990], Asensio Ramos et al. [2007]), materials modeling (Sabin et al. [2000]) and the detection and diagnosis of disease (Siltanen et al. [2003], Kolehmainen et al. [2006]). The two critical ingredients of a Bayesian inference problem are - an informative prior representing the prior belief about the parameters and an efficient method for sampling from the posterior distribution. In this manuscript we describe how a deep generative model (generative adversarial networks (GANs)) can be used in these roles.

In a typical inverse problem, we wish to infer a vector of parameters $\boldsymbol{x} \in \mathbb{R}^N$ from the measurement of a related vector $\boldsymbol{y} \in \mathbb{R}^P$, where the two are related through a forward model $\boldsymbol{y} = \boldsymbol{f}(\boldsymbol{x})$. A noisy measurement of $\boldsymbol{y}$ is denoted by $\hat{\boldsymbol{y}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{\eta}$, where $\boldsymbol{\eta} \in \mathbb{R}^P$ represents noise. While the forward map is typically well-posed, its inverse is not, and hence to infer $\boldsymbol{x}$ from the measurement $\hat{\boldsymbol{y}}$ requires techniques that account for this ill-posedness. Classical techniques based on regularization tackle this ill-posedness by using additional information about the sought parameter field explicitly or implicitly (Tarantola [2005]). Bayesian inference offers a different solution to this problem by modeling the unknown parameter and the measurements as random variables and allows for the characterization of the uncertainty in the inferred parameter field. For additive noise, the posterior distribution of $\boldsymbol{x}$, determined using Bayes' theorem after accounting for the observation $\hat{\boldsymbol{y}}$ is given by

$$p_X^{\text{post}}(\boldsymbol{x}|\boldsymbol{y}) = \frac{1}{\mathbb{Z}} p^{\text{l}}(\boldsymbol{y}|\boldsymbol{x}) p_X^{\text{prior}}(\boldsymbol{x}) = \frac{1}{\mathbb{Z}} p_\eta(\hat{\boldsymbol{y}} - \boldsymbol{f}(\boldsymbol{x})) p_X^{\text{prior}}(\boldsymbol{x}), \tag{1}$$

where $\mathbb{Z}$ is the prior-predictive distribution of $\boldsymbol{y}$, $p_X^{\text{prior}}(\boldsymbol{x})$ is the prior distribution of $\boldsymbol{x}$, and $p^{\text{l}}(\boldsymbol{y}|\boldsymbol{x})$ is the likelihood, often determined by the distribution of the error in the model, denoted by p_η.

First Workshop on Solving Inverse Problems with Deep Networks (NeurIPS 2019), Vancouver, Canada.

Despite its numerous applications, Bayesian inference faces significant challenges. These include constructing a reliable and informative prior distribution from a collection of prior measurements denoted by the $\mathcal{S} = \{\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(S)}\}$, and efficiently sampling from the posterior distribution when the dimension of $\boldsymbol{x}$ is large. In this work we consider the use of GANs (Goodfellow et al. [2014]) in addressing these challenges. These networks are useful in this role because of (a) they are able to generate samples of $\boldsymbol{x}$ from $p_X^{\text{gen}}(\boldsymbol{x})$ while ensuring closeness (in an appropriate measure) between $p_X^{\text{gen}}(\boldsymbol{x})$ and the true distribution, and (b) because they accomplish this by sampling from the much simpler distribution of the latent vector $\boldsymbol{z}$, whose dimension is much smaller than that of $\boldsymbol{x}$.

Related work and our contribution: The main idea in this work involves training a GAN using the sample set $\mathcal{S}$, and then using the distribution learned by the GAN as the prior distribution in Bayesian inference. This leads to a useful method for representing complex prior distributions and an efficient approach for sampling from the posterior distribution in terms of the latent vector $\boldsymbol{z}$.

The solution of inverse problems using sample-based priors has a rich history (see Vauhkonen et al. [1997], Calvetti and Somersalo [2005] for example). As does the idea of dimension reduction in parameter space (Marzouk and Najm [2009], Lieberman et al. [2010]). However, the use of GANs in these tasks is novel. Recently, a number of authors have considered the use machine learning-based methods for solving inverse problems. These include the use of convolutional neural networks (CNNs) to solve physics-driven inverse problems (Adler and Öktem [2017], Jin et al. [2017], Patel et al. [2019]), and GANs to solve problems in computer vision (Chang et al., Kupyn et al. [2018], Yang et al. [2018], Ledig et al., Anirudh et al. [2018], Isola et al. [2016], Zhu et al. [2017], Kim et al. [2017]). There is also a growing body of work on using GANs to learn regularizers in inverse problems (Lunz et al. [2018]) and in compressed sensing (Bora et al. [2017, 2018], Kabkab et al. [2018], Wu et al. [2019], Shah and Hegde [2018]). However, these approaches differ from ours in that they solve the inverse problem as an optimization problem and do not quantify uncertainty in a Bayesian framework . More recently, the approach described in (Adler and Öktem [2018]) utilizes GANs in a Bayesian setting; however the GAN is trained to approximate the posterior distribution, and training is done in a supervised fashion with paired samples of the measurement $\hat{\boldsymbol{y}}$ and the corresponding true solution $\boldsymbol{x}$.

2 Problem formulation

Let $\boldsymbol{z} \sim p_Z(\boldsymbol{z})$ characterize the latent vector space and $\boldsymbol{g}(\boldsymbol{z})$ be the generator of a GAN trained using $\mathcal{S}$. Then with infinite capacity and sufficient data, the generator learns the true distribution (Goodfellow et al. [2014]). That is, $p_X^{\text{gen}}(\boldsymbol{x}) = p_X^{\text{true}}(\boldsymbol{x})$, where the distribution $p_X^{\text{gen}}(\boldsymbol{x})$ is defined as

$$\boldsymbol{x} \sim p_X^{\text{gen}}(\boldsymbol{x}) \Rightarrow \boldsymbol{x} = \boldsymbol{g}(\boldsymbol{z}), \boldsymbol{z} \sim p_Z(\boldsymbol{z}). \tag{2}$$

Here p_Z is the multivariate distribution of the latent vector whose components are iid and typically conform to a Gaussian or a uniform distribution. Now consider a measurement $\hat{\boldsymbol{y}}$ for which we would like to infer the posterior distribution of $\boldsymbol{x}$. For this we use (1) and set the prior distribution to be equal to the true distribution, that is $p_X^{\text{prior}} = p_X^{\text{true}} = p_X^{\text{gen}}(\boldsymbol{x})$. Therefore,

$$p_X^{\text{post}}(\boldsymbol{x}|\boldsymbol{y}) \quad = \quad \frac{1}{\mathbb{Z}} p_\eta(\hat{\boldsymbol{y}} - \boldsymbol{f}(\boldsymbol{x})) p_X^{\text{gen}}(\boldsymbol{x}). \tag{3}$$

Using this it is easy to show that for any $l(\boldsymbol{x})$,

$$\mathop{\mathbb{E}}_{\boldsymbol{x} \sim p_X^{\text{post}}}[l(\boldsymbol{x})] \quad = \quad \mathop{\mathbb{E}}_{\boldsymbol{z} \sim p_Z^{\text{post}}}[l(\boldsymbol{g}(\boldsymbol{z}))], \tag{4}$$

where $\mathbb{E}$ is the expectation operator, and

$$p_Z^{\text{post}}(\boldsymbol{z}|\boldsymbol{y}) \equiv \frac{1}{\mathbb{Z}} p_\eta(\hat{\boldsymbol{y}} - \boldsymbol{f}(\boldsymbol{g}(\boldsymbol{z}))) p_Z(\boldsymbol{z}). \tag{5}$$

Note that the distribution p_Z^{post} is the analog of p_X^{post} in the latent vector space. The measurement $\hat{\boldsymbol{y}}$ updates the prior distribution for $\boldsymbol{x}$ to the posterior distribution; similarly, it updates the prior distribution for $\boldsymbol{z}$, p_Z, to the posterior distribution, p_Z^{post}. Equation (4) implies that sampling from the posterior distribution for $\boldsymbol{x}$ is equivalent to sampling from the posterior distribution for $\boldsymbol{z}$ and transforming the sample through the generator $\boldsymbol{g}$. That is,

$$\boldsymbol{x} \sim p_X^{\text{post}}(\boldsymbol{x}|\boldsymbol{y}) \Rightarrow \boldsymbol{x} = \boldsymbol{g}(\boldsymbol{z}), \boldsymbol{z} \sim p_Z^{\text{post}}(\boldsymbol{z}|\boldsymbol{y}). \tag{6}$$

Since the dimension of $\boldsymbol{z}$ is typically smaller than that of $\boldsymbol{x}$, and since the operation of the generator is typically inexpensive, this represents an efficient approach to sampling from the posterior of $\boldsymbol{x}$.

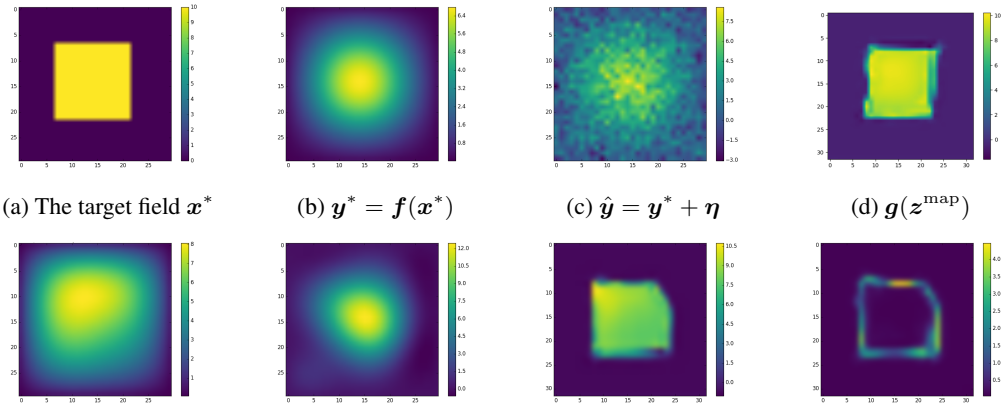

(a) The target field $\boldsymbol{x}^*$ (b) $\boldsymbol{y}^* = \boldsymbol{f}(\boldsymbol{x}^*)$ (c) $\hat{\boldsymbol{y}} = \boldsymbol{y}^* + \boldsymbol{\eta}$ (d) $\boldsymbol{g}(\boldsymbol{z}^{\text{map}})$

(e) MAP estimate with an H^1 prior (f) MAP estimate with an L_2 prior (g) Mean estimate from MCMC (h) Stddev. estimate from MCMC

Figure 1: Statistics of inferred parameter field.

2.1 Sampling from the posterior distribution

As mentioned in section 1, we wish to infer and characterize the uncertainty in the vector of parameters $\boldsymbol{x}$ from a noisy measurement $\hat{\boldsymbol{y}}$, where $\boldsymbol{f}$ is a known map that connects $\boldsymbol{x}$ and $\boldsymbol{y}$. We also have several prior measurements of $\boldsymbol{x}$, contained in the set $\mathcal{S}$. To solve this problem we train a GAN with a generator $\boldsymbol{g}(\boldsymbol{z})$ on $\mathcal{S}$, and then sample $\boldsymbol{x}$ from $p_X^{\text{post}}(\boldsymbol{x}|\boldsymbol{y})$ given in (6). Since GANs can be used to represent complex distributions efficiently, this algorithm provides a means of including complex priors that are defined by samples. It also leads to an efficient approach to sampling from $p_X^{\text{post}}(\boldsymbol{x}|\boldsymbol{y})$ since the dimension of $\boldsymbol{z}$ is typically smaller (10^1 - 10^2) than that of $\boldsymbol{x}$ (10^4 - 10^7). In Appendix A we describe approaches based on Monte-Carlo, Markov-Chain Monte-Carlo and MAP estimation for estimating population parameters of the posterior that make use of this observation.

3 Results

A problem motivated by physics We apply our approach to the problem of determining the initial temperature distribution of a solid from a measurement of its current temperature. The inferred field ($\boldsymbol{x}$) is represented on a 32^2 grid on a square and the forward operator is defined by the solution of the time-dependent heat conduction problem with uniform conductivity. This operator maps the initial temperature to the temperature at time $t = 1$, and its discrete version is generated by approximating the time-dependent linear heat conduction equation using central differences in space and backward difference in time. It is assumed that the initial temperature is zero everywhere except in a rectangular region, and it is parameterized by the horizontal and vertical coordinates of two corners of the rectangular region and the value of the temperature field within it. 50,000 initial temperature fields sampled from this distribution are included in the sample set $\mathcal{S}$ used to train a Wasserstein GAN (WGAN-GP (Gulrajani et al. [2017])) with an 8-dimensional latent space with batch size of 64 and learning rate of 0.0002.

The target field we wish to infer is shown in Figure 1a. This field is passed through the forward map to generate the noise-free and the noisy versions (Gaussian with zero mean and unit variance) of the measured field shown in Figure 1b and 1c.

We apply the algorithms developed in the previous section to probe the posterior distribution. We first use these to determine the MAP estimate for the posterior distribution of the latent vector (denoted by $\boldsymbol{z}^{\text{map}}$). The value of $\boldsymbol{g}(\boldsymbol{z}^{\text{map}})$ is shown in Figure 1d. By comparing this with the true value of the inferred field, shown in Figure 1a, we observe that the MAP estimate is very close to the true value. This agreement is remarkable if we recognize that the ratio of noise to signal is around 30%, and also compare the MAP estimates obtained using an H^1 or an L_2 prior (see Figures 1e and 1f) with the true value.

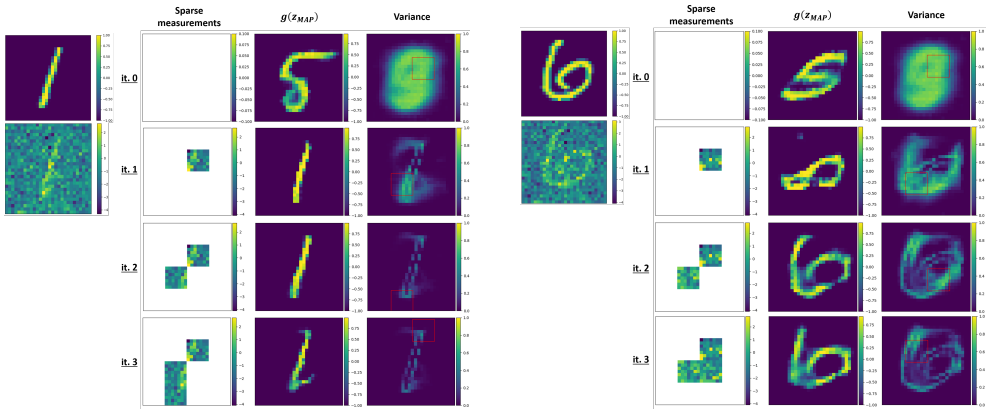

(a) Digit 1 (b) Digit 6

Figure 2: Iterative image recovery with very sparse measurements using uncertainty information: for each digit left most column represents true signal ($\boldsymbol{x}^*$) and its noisy version. The following columns represent the sparse measurement, the estimated MAP, and the estimated variance, respectively at each iteration. The red window in the variance map is the sub-region with maximum variance.

Next, we consider the results obtained by sampling from the MCMC approximation to the posterior distribution of $\boldsymbol{z}$ defined in (5). The MCMC approximation to the mean of the inferred field computed using (8) is shown in Figure 1g. We observe that the edges and the corners of the temperature field are smeared out. This indicates the uncertainty in recovering the values of the initial field along these locations, which can be attributed to the smoothing nature of the forward operator especially for the higher modes. A more precise estimate of the uncertainty in the inferred field is provided by the variance of the inferred initial temperature at each spatial location. In Figure 1h we have plotted the point-wise standard deviation (square-root of the diagonal of co-variance) of the inferred field - our metric of quantified uncertainty. We observe that it is largest along the edges and at the corners, where the forward operator has smoothed out the initial data, and thus introduced large levels of uncertainty in the location of these features. Additional examples of this inverse heat conduction problem with different target fields is shown in Appendix B.

A problem in computer vision: Next we consider a problem in computer vision that highlights the utility of estimating the uncertainty in an inference problem: one of determining the noise-free version of an image from a noisy version of a sub-region of the image. In particular, we consider an iterative version of this problem, where one sub-region is revealed in each iteration, and the user is given the freedom to select this sub-region. We use a strategy that is based on selecting a region where the variance is maximum, and conclude that we arrive at a very good guess for the image in very few iterations. This task falls under active learning regime of machine learning and is useful when measurements are expensive.

We use 55,000 images from the MNIST data set to train a WGAN-GP and use it as a prior in Bayesian inference. We select an image from the complementary set, add Gaussian noise with 0.8 variance, mask regions within this image, and use it to infer the original image. We utilize a forward map that is zero in the masked region and identity everywhere else. We begin by masking the entire image, and allow the user to select the sub-region (which is a square with edge length equal to 1/7th of the original image) in each iteration. We report results when the user selects the sub-region with maximum variance as the sub-region to be revealed in the next iteration. For computing the variance we utilize the algorithm developed in this work.

In Figure 2 we have shown the true image and results from several iterations for two different MNIST digits from test set. For each iteration, we have shown the image that was used as measurement, the corresponding MAP and variance determined using our algorithms. We observe that in the 0^{th} iteration, when nothing is revealed in the measurement, the variance is largest in the center of the image where most digits assume different intensities. This leads to the user requesting a measurement in this region in the subsequent iteration. Thereafter, the estimated variance reduces with each iteration, and we converge to an image which is very close to the true image in very few (2-3) iterations. Additional results for MNIST and CelebA dataset are provided in Appendix B.

References

Jonas Adler and Ozan Öktem. Solving ill-posed inverse problems using iterative deep neural networks. *Inverse Problems*, 33(12):124007, dec 2017. ISSN 0266-5611. doi: 10.1088/1361-6420/aa9581. URL http://stacks.iop.org/0266-5611/33/i=12/a=124007?key=crossref.65c4fa88a47e07d4789aa10592f2090c.

Jonas Adler and Ozan Öktem. Deep bayesian inversion. *arXiv preprint arXiv:1811.05910*, 2018.

Rushil Anirudh, Jayaraman J. Thiagarajan, Bhavya Kailkhura, and Timo Bremer. An Unsupervised Approach to Solving Inverse Problems using Generative Adversarial Networks. may 2018. URL http://arxiv.org/abs/1805.07281.

A. Asensio Ramos, M. J. Martínez González, and J. A. Rubiño-Martín. Bayesian inversion of Stokes profiles. *Astronomy & Astrophysics*, 476(2):959–970, dec 2007. ISSN 0004-6361. doi: 10.1051/0004-6361:20078107. URL http://www.aanda.org/10.1051/0004-6361:20078107.

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.

Ashish Bora, Eric Price, and Alexandros G Dimakis. Ambientgan: Generative models from lossy measurements. *ICLR*, 2:5, 2018.

Daniela Calvetti and Erkki Somersalo. Priorconditioners for linear systems. *Inverse Problems*, 21(4):1397–1418, aug 2005. ISSN 0266-5611. doi: 10.1088/0266-5611/21/4/014. URL http://stacks.iop.org/0266-5611/21/i=4/a=014?key=crossref.ab419ffb66111e3db21bf3d9fd3836f7.

J H Rick Chang, Chun-Liang Li, Barnab Barnabás, P Oczos, B V K Vijaya Kumar, and Aswin C Sankaranarayanan. One Network to Solve Them All-Solving Linear Inverse Problems using Deep Projection Models. Technical report. URL https://arxiv.org/pdf/1703.09912.pdf.

Masoumeh Dashti and Andrew M Stuart. The bayesian approach to inverse problems. *Handbook of Uncertainty Quantification*, pages 1–118, 2016.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

Wences P Gouveia and John A Scales. Resolution of seismic waveform inversion: Bayes versus Occam. *Inverse Problems*, 13(2):323–349, apr 1997. ISSN 0266-5611. doi: 10.1088/0266-5611/13/2/009. URL http://stacks.iop.org/0266-5611/13/i=2/a=009?key=crossref.c3e937d46f2de4adfe9aa2de3c226f3e.

Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.

Tobin Isaac, Noemi Petra, Georg Stadler, and Omar Ghattas. Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet. *Journal of Computational Physics*, 296:348–368, sep 2015. ISSN 0021-9991. doi: 10.1016/J.JCP.2015.04.047. URL https://www.sciencedirect.com/science/article/pii/S0021999115003046.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arxiv*, 2016.

Charles Jackson, Mrinal K. Sen, Paul L. Stoffa, Charles Jackson, Mrinal K. Sen, and Paul L. Stoffa. An Efficient Stochastic Bayesian Approach to Optimal Parameter and Uncertainty Estimation for Climate Model Predictions. *Journal of Climate*, 17(14):2828–2841, jul 2004. ISSN 0894-8755. doi: 10.1175/1520-0442(2004)017<2828:AESBAT>2.0.CO;2. URL http://journals.ametsoc.org/doi/abs/10.1175/1520-0442{%}282004{%}29017{%}3C2828{%}3AAESBAT{%}3E2.0.CO{%}3B2.

Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser. Deep Convolutional Neural Network for Inverse Problems in Imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, sep 2017. ISSN 1057-7149. doi: 10.1109/TIP.2017.2713099. URL http://ieeexplore.ieee.org/document/7949028/.

Maya Kabkab, Pouya Samangouei, and Rama Chellappa. Task-aware compressed sensing with generative adversarial networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160. Springer Science & Business Media, 2006.

Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to Discover Cross-Domain Relations with Generative Adversarial Networks. mar 2017. URL http://arxiv.org/abs/1703.05192.

V. Kolehmainen, A. Vanne, S. Siltanen, S. Jarvenpaa, J.P. Kaipio, M. Lassas, and M. Kalke. Parallelized Bayesian inversion for three-dimensional dental X-ray imaging. *IEEE Transactions on Medical Imaging*, 25(2):218–228, feb 2006. ISSN 0278-0062. doi: 10.1109/TMI.2005.862662. URL http://ieeexplore.ieee.org/document/1583768/.

Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. DeblurGAN: Blind Motion Deblurring Using Conditional Adversarial Networks, 2018. URL http://openaccess.thecvf.com/content{_}cvpr{_}2018/html/Kupyn{_}DeblurGAN{_}Blind{_}Motion{_}CVPR{_}2018{_}paper.html.

Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi Twitter. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. Technical report. URL http://openaccess.thecvf.com/content{_}cvpr{_}2017/papers/Ledig{_}Photo-Realistic{_}Single{_}Image{_}CVPR{_}2017{_}paper.pdf.

Chad Lieberman, Karen Willcox, and Omar Ghattas. Parameter and State Model Reduction for Large-Scale Statistical Inverse Problems. *SIAM Journal on Scientific Computing*, 32(5):2523–2542, jan 2010. ISSN 1064-8275. doi: 10.1137/090775622. URL http://epubs.siam.org/doi/10.1137/090775622.

T. J. Loredo. From Laplace to Supernova SN 1987A: Bayesian Inference in Astrophysics. In *Maximum Entropy and Bayesian Methods*, pages 81–142. Springer Netherlands, Dordrecht, 1990. doi: 10.1007/978-94-009-0683-9_6. URL http://link.springer.com/10.1007/978-94-009-0683-9{_}6.

Sebastian Lunz, Ozan Öktem, and Carola-Bibiane Schönlieb. Adversarial regularizers in inverse problems. In *Advances in Neural Information Processing Systems*, pages 8507–8516, 2018.

James Martin, Lucas C. Wilcox, Carsten Burstedde, and Omar Ghattas. A Stochastic Newton MCMC Method for Large-Scale Statistical Inverse Problems with Application to Seismic Inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, jan 2012. ISSN 1064-8275. doi: 10.1137/110845598. URL http://epubs.siam.org/doi/10.1137/110845598.

Youssef M. Marzouk and Habib N. Najm. Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *Journal of Computational Physics*, 228(6):1862–1902, apr 2009. ISSN 0021-9991. doi: 10.1016/J.JCP.2008.11.024. URL https://www.sciencedirect.com/science/article/pii/S0021999108006062.

H. N. Najm, B. J. Debusschere, Y. M. Marzouk, S. Widmer, and O. P. Le MaÃ®tre. Uncertainty quantification in chemical systems. *International Journal for Numerical Methods in Engineering*, 80(6â7):789–814, nov 2009. ISSN 00295981. doi: 10.1002/nme.2551. URL http://doi.wiley.com/10.1002/nme.2551.

Dhruv Patel, Raghav Tibrewala, Adriana Vega, Li Dong, Nicholas Hugenberg, and Assad A Oberai. Circumventing the solution of inverse problems in mechanics through deep learning: Application to elasticity imaging. *Computer Methods in Applied Mechanics and Engineering*, 353:448–466, 2019.

Adriano Polpo, Julio Stern, Francisco Louzada, Rafael Izbicki, and Hellinton Takada, editors. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 239 of *Springer Proceedings in Mathematics & Statistics*. Springer International Publishing, Cham, 2018. ISBN 978-3-319-91142-7. doi: 10.1007/978-3-319-91143-4. URL http://link.springer.com/10.1007/978-3-319-91143-4.

T J Sabin, C A L Bailer-Jones, and P J Withers. Accelerated learning using Gaussian process models to predict static recrystallization in an Al-Mg alloy. *Modelling and Simulation in Materials Science and Engineering*, 8(5):687–706, sep 2000. ISSN 0965-0393. doi: 10.1088/0965-0393/8/5/304. URL http://stacks.iop.org/0965-0393/8/i=5/a=304?key=crossref.0a7cd40dc84219e2de59f113e710378b.

Viraj Shah and Chinmay Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4609–4613. IEEE, 2018.

S Siltanen, V Kolehmainen, S J rvenp, J P Kaipio, P Koistinen, M Lassas, J Pirttil, and E Somersalo. Statistical inversion for medical x-ray tomography with few radiographs: I. General theory. *Physics in Medicine and Biology*, 48(10):1437–1463, may 2003. ISSN 0031-9155. doi: 10.1088/0031-9155/48/10/314. URL http://stacks.iop.org/0031-9155/48/i=10/a=314?key=crossref.5fce2d21d49cf69f7c7b946fb1945c85.

Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*, volume 89. siam, 2005.

M Vauhkonen, J P Kaipio, E Somersalo, and P A Karjalainen. Electrical impedance tomography with basis constraints. *Inverse Problems*, 13(2):523–530, apr 1997. ISSN 0266-5611. doi: 10.1088/0266-5611/13/2/020. URL http://stacks.iop.org/0266-5611/13/i=2/a=020?key=crossref.46559bf45aab26a8302acc14e8db4c89.

Jingbo Wang and Nicholas Zabaras. Hierarchical bayesian models for inverse problems in heat conduction. *Inverse Problems*, 21(1):183–206, dec 2004. doi: 10.1088/0266-5611/21/1/012. URL https://doi.org/10.1088%2F0266-5611%2F21%2F1%2F012.

Yan Wu, Mihaela Rosca, and Timothy Lillicrap. Deep compressed sensing. *arXiv preprint arXiv:1905.06723*, 2019.

Qingsong Yang, Pingkun Yan, Yanbo Zhang, Hengyong Yu, Yongyi Shi, Xuanqin Mou, Mannudeep K. Kalra, Yi Zhang, Ling Sun, and Ge Wang. Low-Dose CT Image Denoising Using a Generative Adversarial Network With Wasserstein Distance and Perceptual Loss. *IEEE Transactions on Medical Imaging*, 37(6):1348–1357, jun 2018. ISSN 0278-0062. doi: 10.1109/TMI.2018.2827462. URL https://ieeexplore.ieee.org/document/8340157/.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.

Appendix A: Posterior characterization

Monte-Carlo (MC) approximation From (4) we conclude,

$$\overline{l(\boldsymbol{x})} \equiv \mathop{\mathbb{E}}_{\boldsymbol{x}\sim p_X^{\text{post}}}[l(\boldsymbol{x})] \approx \frac{\sum_{n=1}^{N_{\text{samp}}} l(\boldsymbol{g}(\boldsymbol{z}))p_\eta(\hat{\boldsymbol{y}} - \boldsymbol{f}(\boldsymbol{g}(\boldsymbol{z})))}{\sum_{n=1}^{N_{\text{samp}}} p_\eta(\hat{\boldsymbol{y}} - \boldsymbol{f}(\boldsymbol{g}(\boldsymbol{z})))}, \qquad \boldsymbol{z} \sim p_Z(\boldsymbol{z}). \tag{7}$$

The sampling in this approach is rather simple since in a typical GAN the z_is are sampled from simple distributions like a Gaussian or a uniform distribution.

Markov-Chain Monte-Carlo (MCMC) approximation A more efficient approach is to generate an MCMC approximation $p_Z^{\text{mcmc}}(\boldsymbol{z}|\boldsymbol{y}) \approx p_Z^{\text{post}}(\boldsymbol{z}|\boldsymbol{y})$ using (5), and sampling $\boldsymbol{z}$ from this distribution,

$$\overline{l(\boldsymbol{x})} \equiv \mathop{\mathbb{E}}_{\boldsymbol{x}\sim p_X^{\text{post}}}[l(\boldsymbol{x})] \approx \frac{1}{N_{\text{samp}}} \sum_{n=1}^{N_{\text{samp}}} l(\boldsymbol{g}(\boldsymbol{z})), \qquad \boldsymbol{z} \sim p_Z^{\text{mcmc}}(\boldsymbol{z}|\boldsymbol{y}). \tag{8}$$

Expression for the maximum a-posteriori estimate When the components of the latent vector are iid with a normal distribution with zero mean and unit variance and the components of noise vector are multivariate normal with zero mean and a covariance matrix $\boldsymbol{\Sigma}$, from (5), we have

$$p_Z^{\text{post}}(\boldsymbol{z}|\boldsymbol{y}) \propto \exp\Big(-\frac{1}{2}\overbrace{\big(|\boldsymbol{\Sigma}^{-1/2}(\hat{\boldsymbol{y}} - \boldsymbol{f}(\boldsymbol{g}(\boldsymbol{z})))|^2 + |\boldsymbol{z}|^2\big)}^{\equiv r(\boldsymbol{z})}\Big). \tag{9}$$

The MAP estimate for this distribution is $\boldsymbol{z}^{\text{map}} = \arg\min_{\boldsymbol{z}} r(\boldsymbol{z})$. This minimization problem may be solved using any gradient-based optimization algorithm. Once $\boldsymbol{z}^{\text{map}}$ is determined, one may evaluate $\boldsymbol{g}(\boldsymbol{z}^{\text{map}})$ by using the GAN generator.

Appendix B: Additional results

B.1 Additional examples for physics motivated inference problem

Figure 3 shows additional examples of inverse heat conduction problem described in section 3 for different prior distributions. First two rows show results for the case where we consider a family of initial temperatures where the background is zero, and the temperature on a rectangular sub-domain varies linearly from 2 units on the left edge to 4 units on the right edge. In the last two rows we show results for the case where the possible initial temperature distribution (i.e. prior distribution) is MNIST digits.

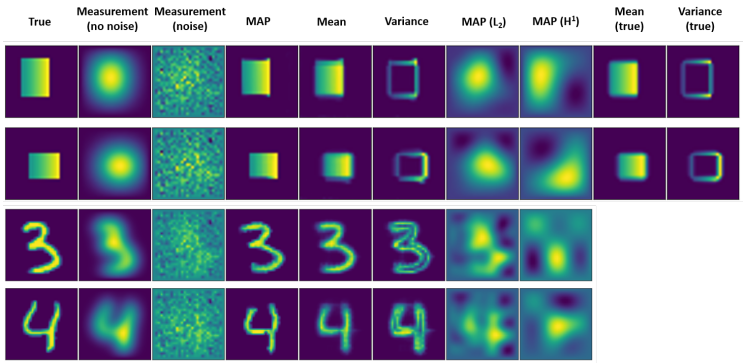

Figure 3: From left to right: (1) true initial temperature, (2) temperature at $t = 1$, (3) noisy version temperature used as measurement, (4), (5), (6) MAP, mean and pixel-wise variance estimates using GAN priors, (7) and (8) MAP estimates using L_2 and H^1 Gaussian priors, (9) & (10) true MAP and variance obtained by sampling over the true parameter space.

B.2 Additional results for active learning applied to computer vision applications

Here we provide additional examples of iterative image recovery scheme described in section 3 for MNIST (figure 4) and CelebA (figure 6) dataset. We also compare the performance of this variance-driven iterative strategy to random sampling scheme, where the next sub-region is selected randomly (figure 5).

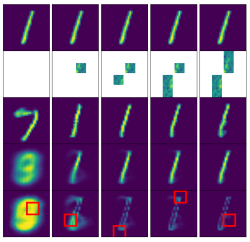

(a) Digit 1

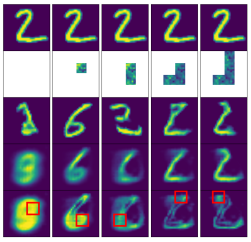

(b) Digit 2

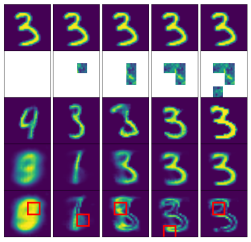

(c) Digit 3

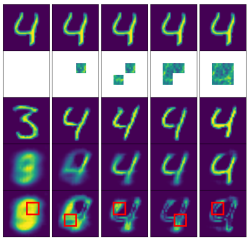

(d) Digit 4

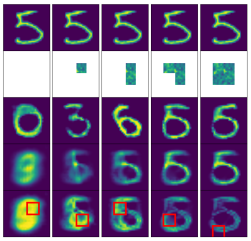

(e) Digit 5

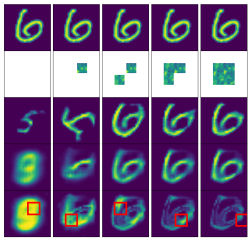

(f) Digit 6

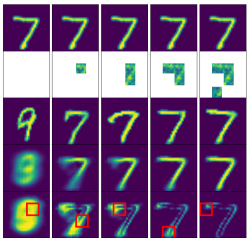

(g) Digit 7

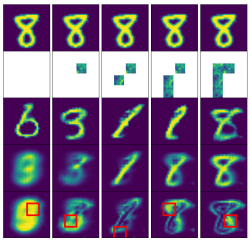

(h) Digit 8

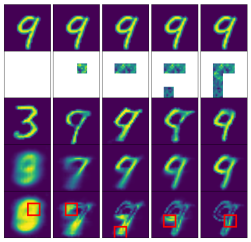

(i) Digit 9

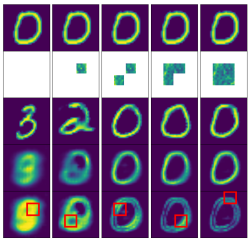

(j) Digit 0

Figure 4: Estimate of the MAP (3rd row), mean (4th row) and variance (5th row) from the limited view of a noisy image (2nd row) using the proposed method. The window to be revealed at a given iteration (shown in red box) is selected using a variance-driven strategy. Top row indicates ground truth. For all digits additive Gaussian noise with variance=0.8 is used.

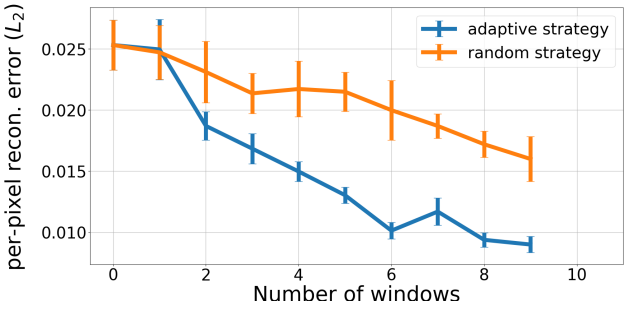

Figure 5: Comparison of variance-driven adaptive strategy and random strategy for MNIST dataset.

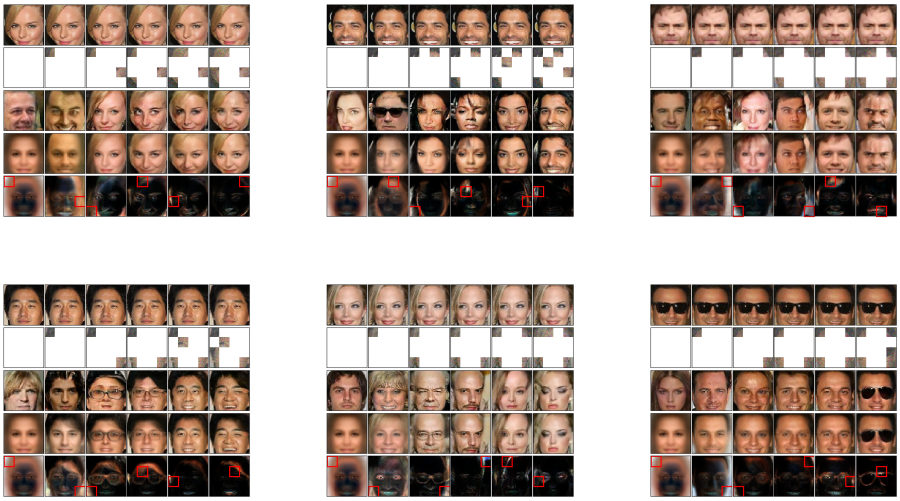

Figure 6: Estimate of the MAP (3rd row), mean (4th row) and variance (5th row) from the limited view of a noisy image (2nd row) using the proposed method. The window to be revealed at a given iteration (shown in red box) is selected using a variance-driven strategy. Top row indicates ground truth. For all images additive Gaussian noise with variance=1 is used.