Variational Generative Flows for Reconstruction Uncertainty Estimation

Jiaxin Zhang [1] **Jan Drgona** [2] **Sayak Mukherjee** [2] **Mahantesh Halappanavar** [2] **Frank Liu** [1]

Abstract

The goal of inverse learning is to determine hidden information from a set of observed but partial measurements. To fully characterize the uncertainty naturally induced by the partial view, a robust inverse solver that is able to estimate the complete posterior of the unrecoverable targets conditioned on a specific observation is therefore important, with a potential to probabilistically interpret the observational data for decision making. In this work, we propose an efficient variational approach that leverages a generative model to learn an approximate posterior distribution for the purpose of quantifying uncertainty in hidden targets. This is achieved by parameterizing the target posterior using a flow-based model and minimizing the KL divergence between the generative distribution and the posterior distribution. Without requiring large training data, the target posterior samples can be efficiently drawn from the learned flow-based model through an invertible transformation from tractable Gaussian random samples. We demonstrate our proposed approach on a real-world FastMRI image reconstruction problem and find it achieves high-quality performance with a smaller variation and error compared to the state-of-the-art baseline methods.

1. Introduction

In computer vision and image processing, computational image reconstruction is a typical inverse problem where the goal is to learn and recover a hidden image $\mathbf{x}$ from directly measured data $\mathbf{y}$ via a forward operator $\mathcal{F}$. Such mapping $\mathbf{y} = \mathcal{F}(\mathbf{x})$, referred to as the *forward process* is often well-established. Unfortunately, the *inverse process* $\mathbf{x} = \mathcal{F}^{-1}(\mathbf{y})$, proceeds in the opposite direction, which is a nontrivial task since it is often ill-posed. A regularized optimization is therefore formulated to recover the hidden image $\mathbf{x}^*$:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \{\mathcal{L}(\mathbf{y}, \mathcal{F}(\mathbf{x})) + \lambda\, \omega(\mathbf{x})\} \tag{1}$$

where $\mathcal{L}$ is a loss function to measure the difference between the observed measurement data and the forward prediction, ω is a regularization function and λ is a regularization coefficient. The regularization function, including ℓ_1-norm and total variation (VT), are typically used to constrain the image to a unique inverse solution in underdetermined imaging systems (Bouman & Sauer, 1993; Strong & Chan, 2003).

Recent trends have focused on using deep learning for computational image reconstruction, which does not rely on an explicit transformer model or iterative updates but performs learned inversion from representative large datasets (Zhu et al., 2018; Belthangady & Royer, 2019; Tonolini et al., 2020; Wang et al., 2020), with applications in medical science, biology, astronomy and more. However, most of these existing studies in regularized optimization (Natterer & Wübbeling, 2001; Park et al., 2003) and feed-forward deep learning approaches (Ulyanov et al., 2018; Belthangady & Royer, 2019; Wang et al., 2020) mainly focus on pursuing a unique inverse solution by recovering a single point estimate. This leads to a significant limitation when working with underdetermined systems where it is conceivable that multiple inverse image solutions would be equally consistent with the measured data (Barbano et al., 2020; Sun & Bouman, 2020). Practically, in many cases, only partial and limited measurements are available which naturally leads to a *reconstruction uncertainty*. Thus, a reconstruction using a point estimate without uncertainty quantification would potentially mislead the decision making process (Beliy et al., 2019; Zhang et al., 2019; Zhou et al., 2020). Therefore, the ability to characterize and quantify reconstruction uncertainty is of paramount relevance. In principle, Bayesian methods are an attractive route to address the inverse problems with uncertainty estimation. However, in practice, exact Bayesian treatment of complex real-world problems is usually intractable. The common limitation is to resort to inference and sampling, typically by Markov Chain Monte Carlo (MCMC), which are often prohibitively expensive for imaging problems due to the *curse of dimensionality*.

Objective This work aims to achieve a reliable image reconstruction with an accurate estimation of *data* uncertainty

[1]Oak Ridge National Laboratory, USA [2]Pacific Northwest National Laboratory, USA. Correspondence to: Jiaxin Zhang <zhangj@ornl.gov>.

Presented at the ICML 2021 Workshop on Uncertainty and Robustness in Deep Learning.,

resulting from measurement noise and sparsity. A suitable flow-based variational approach is proposed to approximate a posterior distribution of an unobserved (target) image.

Contributions We propose an uncertainty-aware framework that leverages a deep variational approach with robust generative flows to address these challenges. Our goal is to perform accurate characterization and quantification of reconstruction uncertainty (data uncertainty) which is due to sparse and noise measurements. We therefore minimize the model uncertainties caused by invertible architectures by introducing a robust flow-based model. We demonstrate our method on recently introduced problems of FastMRI reconstruction and show that it achieves a reliable and high-quality reconstruction with accurate uncertainty estimation.

2. Background

Generative flow-based models Generative models, such as GANs and VAEs, are intractable for explicitly learning the probability density function which plays a fundamental role on uncertainty estimation. Flow-based generative models overcome this difficulty with the help of *normalizing flows* (NFs), which describe the transformation from a latent density $\mathbf{z}_0 \sim \pi_0(\mathbf{z}_0)$ to a target density $\tau(\mathbf{x})$, where $\mathbf{x} = \mathbf{z}_K \sim \pi_K(\mathbf{z}_K)$ through a sequence of invertible mappings $\mathcal{T}_k : \mathbb{R}^d \to \mathbb{R}^d, k = 1, ..., K$. By using the change of variables rule

$$\begin{aligned} \tau(\mathbf{x}) = \pi_k(\mathbf{z}_k) &= \pi_{k-1}(\mathbf{z}_{k-1}) \left| \det \frac{\partial \mathcal{T}_k^{-1}}{\partial \mathbf{z}_{k-1}} \right| \\ &= \pi_{k-1}(\mathbf{z}_{k-1}) \left| \det \frac{\partial \mathcal{T}_k}{\partial \mathbf{z}_{k-1}} \right|^{-1}, \end{aligned} \tag{2}$$

the target density $\pi_K(\mathbf{z}_K)$ obtained by successively transforming a random variable $\mathbf{z}_0$ through a chain of K transformations $\mathbf{z}_K = \mathcal{T}_K \circ \cdots \circ \mathcal{T}_1(\mathbf{z}_0)$ is

$$\log \tau(\mathbf{x}) = \log \pi_K(\mathbf{z}_K) = \log \pi_0(\mathbf{z}_0) - \sum_{k=1}^{K} \log \left| \det \frac{\partial \mathcal{T}_k}{\partial \mathbf{z}_{k-1}} \right|,$$

where each transformation $\mathcal{T}_k$ must be sufficiently *expressive* while being theoretically *invertible* and efficient to compute the Jacobian determinant. Affine coupling functions (Dinh et al., 2016; Kingma & Dhariwal, 2018) are often used because they are simple and efficient to compute. However, these benefits come at the cost of expressivity and flexibility; many flows must be stacked to learn complex representation, as shown in Figure 1.

Density Estimation Assuming that samples $\{\boldsymbol{x}_i\}_{i=1}^{M}$ drawn from a probability density $p(\mathbf{x})$ are available, our goal is to learn a flow-based model $\tau_\phi(\mathbf{x})$ parameterized by the vector ϕ by through a transformation $\mathbf{x} = \mathcal{T}(\mathbf{z})$ of a latent density $\pi_0(\mathbf{z})$ with $\mathcal{T} = \mathcal{T}_K \circ \cdots \circ \mathcal{T}_1$ as a K-step flow. This is achieved by minimizing the KL-divergence $\mathrm{D}_{\mathrm{KL}} = \mathrm{KL}(p(\mathbf{x}) \parallel \tau_\phi(\mathbf{x}))$, which is equivalent to maximum likelihood estimation.

Variational Inference The goal is to approximate the posterior distribution p through a variational distribution π_K encoded by a flow-based model $\tau_\phi(\mathbf{x})$, which is tractable to compute and draw samples. This is achieved by minimizing the KL-divergence $D_{\mathrm{KL}} = \mathrm{KL}(\pi_K \parallel p)$, which is equivalent to maximizing an evidence lower bound (ELBO).

Evaluation metrics for generative models Designing indicative evaluation metrics for generative models and samples remains a challenge. A widely used metric for measuring the similarity between real and generated images has been the Fréchet Inception Distance (FID) score (Heusel et al., 2017) but it fails to separate two critical aspects of the quality of generative models: *fidelity* that refers to the degree to which the generated samples resemble the real ones, and *diversity*, which measures whether the generated samples cover the full variability of the real samples. We introduce reliable metrics (*density* and *coverage*) to evaluate the quality of the generated samples and measure the difference from the ground truth. They are defined as

$$\begin{aligned} \text{density} &:= \frac{1}{kM} \sum_{j=1}^{M} \sum_{i=1}^{N} 1_{X_j^G \in \mathcal{B}(X_i, \mathrm{NND}_k(X_i))}, \\ \text{coverage} &:= \frac{1}{N} \sum_{i=1}^{N} 1_{\exists j \text{s.t.} X_j^G \in \mathcal{B}(X_i, \mathrm{NND}_k(X_i))} \end{aligned} \tag{3}$$

where N and M are the number of true and generative samples, $\mathcal{B}(x, r)$ is the sphere around x with radius r, and $\mathrm{NND}_k(X_i)$ denotes the distance from X_i to the k^{th} nearest neighbour (Naeem et al., 2020).

3. Method

Our goal is to build a deep variational framework to accurately estimate the data uncertainty quantified by an approximation of the posterior distribution. The regularized optimization in Eq. (1) can be further written in terms of data fidelity (data fitting loss) and regularity:

$$\begin{aligned} \mathbf{x}^* &= \arg\min_{\mathbf{x}} \{\mathcal{L}_{\mathrm{D}}(\mathbf{y}, \mathcal{F}(\mathbf{x})) + \lambda\omega(\mathbf{x})\} \\ &= \arg\min_{\mathbf{x}} \{ \underbrace{\parallel \mathbf{y} - \mathcal{F}(\mathbf{x}) \parallel^2}_{\text{Data fidelity}} + \underbrace{\lambda\omega(\mathbf{x})}_{\text{Regularity}} \} \end{aligned} \tag{4}$$

Assuming the forward operator $\mathcal{F}$ is known and the measurement noise statistics are given, we can reformulate the inverse problem in a probabilistic way. In Bayesian perspective, the regularized inverse problem in Eq. (4) can be interpreted as a Bayesian inference problem but aims to maximize the posterior distribution by searching a point estimator $\mathbf{x}^*$:

$$\begin{aligned}\mathbf{x}^* &= \arg\max_{\mathbf{x}} \{ \underbrace{\log p(\mathbf{x}|\mathbf{y})}_{\text{Posterior}} \} \\ &= \arg\max_{\mathbf{x}} \{ \underbrace{\log p(\mathbf{y}|\mathbf{x})}_{\text{Data likelihood}} + \underbrace{\log p(\mathbf{x})}_{\text{Prior}} \}\end{aligned} \tag{5}$$

where the prior distribution $p(\mathbf{x})$ (e.g., image prior (Ulyanov et al., 2018) in reconstruction problems) defines a similar regularization term and data likelihood $p(\mathbf{y}|\mathbf{x})$ corresponds to the data fidelity in Eq. (4).

If we parameterize the target $\mathbf{x}$ using a generative model $\mathbf{x} = \mathcal{T}_{\phi}(\mathbf{z}), \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$ with model parameter ϕ, an approximate posterior distribution $\tau_{\phi^*}(\mathbf{x})$ is obtained by minimizing the KL-divergence between the generative distribution and the target posterior distribution

$$\begin{aligned}\boldsymbol{\phi}^* &= \arg\min_{\boldsymbol{\phi}} \mathrm{KL}(\tau_{\boldsymbol{\phi}}(\mathbf{x}) \parallel p(\mathbf{x}|\mathbf{y})) \\ &= \arg\min_{\boldsymbol{\phi}} \mathbb{E}_{\mathbf{x}\sim\tau_{\boldsymbol{\phi}}(\mathbf{x})}[-\log p(\mathbf{y}|\mathbf{x}) - \log p(\mathbf{x}) + \log \tau_{\boldsymbol{\phi}}(\mathbf{x})]\end{aligned}$$

Unfortunately, the probability density (data likelihood) $\tau_{\phi}(\mathbf{x})$ can not be exactly evaluated by most of existing generative models, such as GANs (Goodfellow et al., 2014) or VAEs (Kingma & Welling, 2013). Flow-based models (Rezende & Mohamed, 2015; Dinh et al., 2016; Kingma & Dhariwal, 2018; Grathwohl et al., 2018; Wu et al., 2020; Nielsen et al., 2020) offer a promising approach to compute the likelihood exactly via the change of variable theorem with invertible architectures. Therefore, Eq. (3) can be reformulated in terms of flow-based model as

$$\begin{aligned}\boldsymbol{\phi}^* = \arg\min_{\boldsymbol{\phi}} \mathbb{E}_{\mathbf{z}\sim\pi(\mathbf{z})} [&-\log p(\mathbf{y}|\mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z})) - \log p(\mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z})) \\ &+ \log \pi(\mathbf{z}) - \log \left|\det \nabla_{\mathbf{z}} \mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z})\right|\end{aligned}$$

We replace data likelihood and prior terms by using data fidelity loss and regularization function in Eq. (4) such that we can define a new optimization problem where it can be solved by using a Monte Carlo method to approximate the expectation in practice:

$$\begin{aligned}\boldsymbol{\phi}^* &= \arg\min_{\boldsymbol{\phi}} \mathbb{E}_{\mathbf{z}\sim\pi(\mathbf{z})} [\mathcal{L}_{\mathrm{D}}(\mathbf{y}, \mathcal{F}(\mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z}))) \\ &\quad + \lambda\omega(T_{\boldsymbol{\phi}}(\mathbf{z})) + \log \pi(\mathbf{z}) - \log \left|\det \nabla_{\mathbf{z}} \mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z})\right| \\ &= \arg\min_{\boldsymbol{\phi}} \sum_{j=1}^{M} \Big[\mathcal{L}_{\mathrm{D}}(\mathbf{y}, \mathcal{F}(\mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z}_j))) + \lambda\omega(\mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z}_j)) \\ &\quad - \underbrace{\log \left|\det \nabla_{\mathbf{z}} \mathcal{T}_{\boldsymbol{\phi}}(\mathbf{z}_j)\right|}_{\text{Entropy}} \Big]\end{aligned}$$

where $\pi(\mathbf{z})$ is a constant under expectation and $\log \left|\det \nabla_{\mathbf{z}} \mathcal{T}_{\phi}(\mathbf{z}_j)\right|$ is an entropy term that is important to encourage sample diversity and exploration so as to avoid generative model from collapsing to a deterministic solution (Higgins et al., 2016).

It can be noted that the flow-based model is very critical and sensitive to uncertainty estimation and quantification within this variational framework. To perform accurate data uncertainty estimation, the uncertainty associated with the flow-based model must be minimized. To this end, we propose to use a robust generative flow (RGF) by leveraging neural spline flow-based models (Durkan et al., 2019) with enhanced stability, expressivity and flexibility, while conserving efficient inference and sampling without increasing architecture depth.

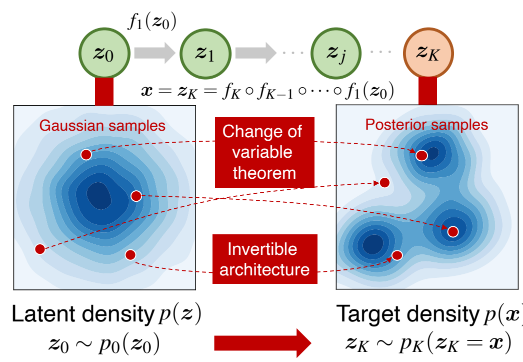

Figure 1: Illustration of mechanism for normalizing flows. The prior Gaussian samples drawn from the latent distribution is transformed to the posterior samples that match the target density by using a sequence of invertible mappings.

4. Experiments

FastMRI case study Partial and undersampled noisy measurements in MRI will lead to reconstruction uncertainty. We demonstrate that our proposed variational framework with robust generative flows (RGF) and variance-reduced sampling (LPSS) (Shields & Zhang, 2016) can be successfully applied to quantify the reconstruction (data) uncertainty and error on two cases (brain and knee) from FastMRI dataset (Zbontar et al., 2018) (resized to 128 × 128 pixels) with three different acceleration factors 4X, 6X and 8X. Fig. 2 presents the reconstruction results with pixel-wise statistics of the estimated posterior distribution. 1000 posterior samples drawn from the learned models are used to estimate the statistical information. Note that, for both brain and knee cases with a speedup 4X factor, our RGFL shows a more accurate mean estimate $\mu_{\mathbf{x}}$ with a smaller absolute error $\epsilon_{\mathbf{x}}$ than the baseline RealNVP model with similar architecture (4 blocks). Our advantage in terms of standard deviation $\sigma_{\mathbf{x}}$ is more significant thanks to the model robustness with variance reduction. In other words, our method provides a more reliable reconstruction given the same measurement. As expected, the pixel-wise $\sigma_{\mathbf{x}}$ of the reconstruction tends to be larger as the speedup factor increases (more details in supplementary materials).

Table 1: Statistical comparison of the estimated posteriors on FastMRI brain case

Brain	Speedup 4X			Speedup 6X			Speedup 8X		
	RNVP	NSF	RGFL	RNVP	NSF	RGFL	RNVP	NSF	RGFL
Std. Dev. ↓	1.21E-5	8.82E-6	**9.73E-7**	1.97E-5	1.02E-5	**2.35E-6**	3.83E-5	2.27E-5	**4.10E-6**
Abs. Error ↓	3.31E-5	6.91E-6	**1.78E-6**	5.70E-5	8.04E-6	**4.58E-6**	8.81E-5	1.33E-5	**7.12E-6**
Precision ↑	0.548	0.564	**0.603**	0.501	0.555	**0.590**	0.511	0.540	**0.566**
Recall ↑	0.564	0.598	**0.629**	0.523	0.582	**0.606**	0.493	**0.571**	0.569
Density ↑	0.925	0.934	**0.955**	0.887	0.929	**0.953**	0.831	0.907	**0.940**
Coverage ↑	0.957	0.989	**0.997**	0.892	0.983	**0.988**	0.887	0.965	**0.971**

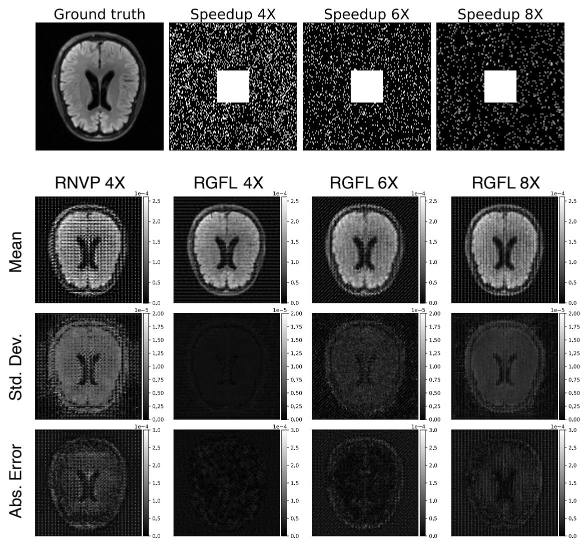

Figure 2: FastMRI reconstruction of brain case at three different acceleration speedup factors: 4X, 6X and 8X (each shown in a column). Row 1 shows the ground truth and sampling masks for each case. Row 2-4 shows the mean, standard deviation, absolution of error for the estimated posterior samples.

We further use the mean of the standard deviation $\overline{\sigma}_{\mathbf{x}}$ and the absolute error $\overline{\epsilon}_{\mathbf{x}}$ to quantitatively compare the pixel-wise statistics (see Table 1). Our method outperforms the other two baselines (RNVP and NSF) in terms of accuracy and variation of the reconstruction. Specifically, our estimation achieves significant variance reduction with 1 or 2 orders of magnitude. The fidelity and diversity metrics are used here to evaluate the posterior samples drawn from the learned generative models (↓ and ↑ means a lower and higher value is expected respectively). Our method shows competitive performance in most cases.

5. Related works

Deep learning for solving inverse problems requires uncertainty estimation to be reliable in real settings. Bayesian deep learning (Kendall & Gal, 2017; Khan et al., 2018; Wilson & Izmailov, 2020), specifically Bayesian neural networks (Hernández-Lobato & Adams, 2015; Gal) can achieve this goal while offering a computationally tractable way for recovering reconstruction uncertainty. However, exact inference in the BNN framework is not a trivial task, so several variational approximation approaches are proposed to deal with the scalability challenges. Monte Carlo dropout (Gal & Ghahramani, 2016) can be seen as a promising alternative approach that is easy to implement and evaluate. Deep ensemble (Lakshminarayanan et al., 2017) methods proposed by combining multiple deep models from different initialization have outperformed BNN. Recent methods on deterministic uncertainty quantification (Van Amersfoort et al., 2020; van Amersfoort et al., 2021) use a single forward pass but scales well to large datasets. Although these approaches show impressive performance, they rely on supervised learning with paired input-output datasets and only characterize the uncertainty conditioned on a training set.

Variational methods offer a more efficient alternative approximating true but intractable posterior distribution by an optimally selected tractable distribution family (Blei et al., 2017). However, the restriction to limited distribution families fails if the true posterior is too complex. Recent advances in conditional generative models, such as conditional GANs (cGANs) (Wang et al., 2018), overcome this restriction in principle, but have limitations in satisfactory diversity in practice. Another commonly adopted option is conditional VAEs (cVAEs) (Sohn et al., 2015), which outperform cGANs in some cases, but in fact, the direct application of both conditional generative models in computational imaging is challenging because a large number of data is typically required (Tonolini et al., 2020). This introduces additional difficulties if our observations and measurements are sparse, and expensive to collect.

6. Conclusion

In this work, we propose an uncertainty-aware framework that leverages a deep variational approach with robust generative flows and variance-reduced sampling to perform an accurate estimation of reconstruction uncertainty. We minimize the model uncertainties by developing a robust flow-based model and decrease the sampling variation via variance-reduced sampling. The results on a real-world MRI demonstrate our advantages. The future work focuses on improving the invertiblility of normalizing flows and reducing the computational cost in the variance-reduced sampling.

References

Barbano, R., Kereta, Ž., Zhang, C., Hauptmann, A., Arridge, S., and Jin, B. Quantifying sources of uncertainty in deep learning-based image reconstruction. *arXiv preprint arXiv:2011.08413*, 2020.

Beliy, R., Gaziv, G., Hoogi, A., Strappini, F., Golan, T., and Irani, M. From voxels to pixels and back: Self-supervision in natural-image reconstruction from fmri. *arXiv preprint arXiv:1907.02431*, 2019.

Belthangady, C. and Royer, L. A. Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction. *Nature methods*, 16(12):1215–1225, 2019.

Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.

Bouman, C. and Sauer, K. A generalized gaussian image model for edge-preserving map estimation. *IEEE Transactions on image processing*, 2(3):296–310, 1993.

Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Neural spline flows. *arXiv preprint arXiv:1906.04032*, 2019.

Gal, Y. Uncertainty in deep learning.

Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.

Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.

Hernández-Lobato, J. M. and Adams, R. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, pp. 1861–1869. PMLR, 2015.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6629–6640, 2017.

Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.

Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, 2017.

Khan, M., Nielsen, D., Tangkaratt, V., Lin, W., Gal, Y., and Srivastava, A. Fast and scalable bayesian deep learning by weight-perturbation in adam. In *International Conference on Machine Learning*, pp. 2611–2620. PMLR, 2018.

Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.

Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NIPS*, 2017.

Naeem, M. F., Oh, S. J., Uh, Y., Choi, Y., and Yoo, J. Reliable fidelity and diversity metrics for generative models. In *International Conference on Machine Learning*, pp. 7176–7185. PMLR, 2020.

Natterer, F. and Wübbeling, F. *Mathematical methods in image reconstruction*. SIAM, 2001.

Nielsen, D., Jaini, P., Hoogeboom, E., Winther, O., and Welling, M. Survae flows: Surjections to bridge the gap between vaes and flows. *Advances in Neural Information Processing Systems*, 33, 2020.

Park, S. C., Park, M. K., and Kang, M. G. Super-resolution image reconstruction: a technical overview. *IEEE signal processing magazine*, 20(3):21–36, 2003.

Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538. PMLR, 2015.

Shields, M. D. and Zhang, J. The generalization of latin hypercube sampling. *Reliability Engineering & System Safety*, 148:96–108, 2016.

Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28: 3483–3491, 2015.

Strong, D. and Chan, T. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse problems*, 19(6):S165, 2003.

Sun, H. and Bouman, K. L. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. *arXiv preprint arXiv:2010.14462*, 2020.

Tonolini, F., Radford, J., Turpin, A., Faccio, D., and Murray-Smith, R. Variational inference for computational imaging inverse problems. *Journal of Machine Learning Research*, 21(179):1–46, 2020.

Ulyanov, D., Vedaldi, A., and Lempitsky, V. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9446–9454, 2018.

Van Amersfoort, J., Smith, L., Teh, Y. W., and Gal, Y. Uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, pp. 9690–9700. PMLR, 2020.

van Amersfoort, J., Smith, L., Jesson, A., Key, O., and Gal, Y. Improving deterministic uncertainty estimation in deep learning for classification and regression. *arXiv preprint arXiv:2102.11409*, 2021.

Wang, G., Ye, J. C., and De Man, B. Deep learning for tomographic image reconstruction. *Nature Machine Intelligence*, 2(12):737–748, 2020.

Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., Tao, A., Kautz, J., and Catanzaro, B. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8798–8807, 2018.

Wilson, A. G. and Izmailov, P. Bayesian deep learning and a probabilistic perspective of generalization. *arXiv preprint arXiv:2002.08791*, 2020.

Wu, H., Köhler, J., and Noé, F. Stochastic normalizing flows. *arXiv preprint arXiv:2002.06707*, 2020.

Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M. J., Defazio, A., Stern, R., Johnson, P., Bruno, M., et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.

Zhang, Z., Romero, A., Muckley, M. J., Vincent, P., Yang, L., and Drozdzal, M. Reducing uncertainty in undersampled mri reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2049–2058, 2019.

Zhou, Q., Yu, T., Zhang, X., and Li, J. Bayesian inference and uncertainty quantification for medical image reconstruction with poisson data. *SIAM Journal on Imaging Sciences*, 13(1):29–52, 2020.

Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R., and Rosen, M. S. Image reconstruction by domain-transform manifold learning. *Nature*, 555(7697):487–492, 2018.