

# Uncertainty Quantification for Deep Unrolling-Based Computational Imaging

Canberk Ekmekci<sup>1</sup>, *Graduate Student Member, IEEE*, and Mujdat Cetin<sup>1</sup>, *Fellow, IEEE*

**Abstract**—Deep unrolling is an emerging deep learning-based image reconstruction methodology that bridges the gap between model-based and purely deep learning-based image reconstruction methods. Although deep unrolling methods achieve state-of-the-art performance for imaging problems and allow the incorporation of the observation model into the reconstruction process, they do not provide any uncertainty information about the reconstructed image, which severely limits their use in practice, especially for safety-critical imaging applications. In this article, we propose a learning-based image reconstruction framework that incorporates the observation model into the reconstruction task and that is capable of quantifying epistemic and aleatoric uncertainties, based on deep unrolling and Bayesian neural networks. We demonstrate the uncertainty characterization capability of the proposed framework on magnetic resonance imaging and computed tomography reconstruction problems. We investigate the characteristics of the epistemic and aleatoric uncertainty information provided by the proposed framework to motivate future research on utilizing uncertainty information to develop more accurate, robust, trustworthy, uncertainty-aware, learning-based image reconstruction and analysis methods for imaging problems. We show that the proposed framework can provide uncertainty information while achieving comparable reconstruction performance to state-of-the-art deep unrolling methods.

**Index Terms**—Image reconstruction, uncertainty quantification, uncertainty characterization, deep unrolling, computational imaging, learning-based imaging.

## I. INTRODUCTION

THIS article concerns imaging problems where the target image is observed through a linear transformation followed by additive noise. This observation model is quite general and has been used to model a variety of imaging techniques such as computed tomography (CT) [1], magnetic resonance imaging (MRI) [2], microscopy [3], and radar imaging [4].

For this observation model, classical model-based iterative reconstruction methods cast the image reconstruction problem

Manuscript received 6 July 2022; revised 18 November 2022; accepted 17 December 2022. Date of publication 30 December 2022; date of current version 6 January 2023. This work was supported by the National Science Foundation under Grant CCF-1934962. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Lei Tian. (Corresponding author: Canberk Ekmekci.)

Canberk Ekmekci is with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627 USA (e-mail: cekmekci@ur.rochester.edu).

Mujdat Cetin is with the Department of Electrical and Computer Engineering and Goergen Institute for Data Science, University of Rochester, Rochester, NY 14627 USA (e-mail: mujdat.cetin@rochester.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2022.3233185>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2022.3233185

as a regularized least squares problem whose objective function is the sum of a data-fidelity term and a regularizer. The observation model of the imaging problem determines the form of the data-fidelity term, and the prior knowledge about the target image governs the form of the regularizer. After obtaining an analytical expression for the data-fidelity term and choosing a regularization function, such as the total variation (TV) semi-norm [5], the resulting optimization problem is solved iteratively by using an appropriate iterative optimization algorithm such as alternating direction method of multipliers (ADMM) [6], half-quadratic splitting (HQS), and proximal gradient descent (PGD) method [7].

Inspired by model-based iterative reconstruction methods and the pioneering work of Gregor and LeCun [8] on sparse coding, a deep learning-based image reconstruction methodology, which is often referred to as deep unrolling [9], [10], [11], [12], [13], [14], [15], [16] has emerged to bridge the gap between model-based image reconstruction methods and purely deep neural network-based image reconstruction methods. The common theme among deep unrolling methods is that they often design a deep neural network by replacing some parts of the unrolled iterative reconstruction algorithm with trainable parameters and neural networks. The main advantages of deep unrolling methods are that they explicitly incorporate the observation model into the neural network, hence they enforce data consistency, and the resulting deep learning-based image reconstruction method is essentially classical model-based reconstruction algorithm with some learnable components.

Although deep unrolling methods have the advantage of incorporating domain knowledge and the physics of the imaging problem into the neural network architecture, existing deep unrolling methods do not provide any predictive uncertainty information about the reconstructed image since they rely on non-Bayesian (standard) neural networks to reconstruct the target image from the corrupted measurements. This severely limits their applicability in safety-critical real-world imaging applications such as medical imaging, where uncertainty information is crucial to make accurate decisions. Our perspective is that we can solve this problem by taking a Bayesian approach for uncertainty estimation and using Bayesian neural networks (BNNs) [17]. BNNs are probabilistic models that can quantify the inherent uncertainty on the target image for a given measurement vector due to the ill-posed nature of the inverse problem, which is referred to as the *aleatoric* uncertainty [18], and the uncertainty on the parameters of a neural network, which is referred to as the

*epistemic* uncertainty [18], by putting a probability distribution on the parameters and computing the posterior distribution of the parameters given a training dataset. By using BNNs together with the idea of deep unrolling, we claim that we can provide predictive uncertainty information for the reconstructed image while preserving the advantages of deep unrolling.

The contribution of this article is three-fold:

- 1) By bringing the idea of deep unrolling and Bayesian neural networks together, we propose an uncertainty-quantifying learning-based image reconstruction framework. Our approach characterizes the overall predictive uncertainty, which is composed of aleatoric and epistemic uncertainties. The proposed method first assumes that the aleatoric uncertainty has the form of an additive Gaussian noise, which is implicitly assumed by most state-of-the-art deep unrolled networks as shown in Section III-B, and defines a likelihood function such that the aleatoric uncertainty is modeled with a U-shaped neural network [19] and, the mean of the likelihood function is represented with a deep unrolled neural network. Then, by following the principles of Bayesian neural networks, the proposed method approximates the posterior distribution of the parameters of the likelihood function using a scalable variational inference method called Monte Carlo (MC) Dropout [20]. Next, for a given test measurement and a training dataset, the proposed method computes the predictive distribution, which can be later used to obtain the reconstructed image and the epistemic and aleatoric uncertainty maps, via Monte Carlo integration.
- 2) We qualitatively evaluate the proposed method on MRI and CT reconstruction problems to validate whether the epistemic and aleatoric uncertainty information provided by the proposed method exhibits the theoretical characteristics of epistemic and aleatoric uncertainties. We show that the epistemic and aleatoric uncertainty maps obtained by the proposed method display some of the key theoretical characteristics of epistemic and aleatoric uncertainties.
- 3) To assess the quality of the probabilistic predictions quantitatively, we generate the calibration plots and calculate the calibration metrics of the proposed method for MRI and CT reconstruction problems. We show that due to the modeling assumptions made by the proposed method, the proposed method may output slightly underconfident predictions. We later show that the proposed model can be easily calibrated by using the recalibration method introduced by Kuleshov et al. [21] to output calibrated probabilistic predictions.

We show that the proposed framework can achieve comparable reconstruction performance to a state-of-the-art deep unrolling method and provide epistemic and aleatoric uncertainty information about the reconstructed image while incorporating the domain knowledge into the reconstruction process.

## II. RELATED WORK

The problem of uncertainty quantification for image reconstruction tasks, e.g., [22], [23], [24], [25], [26], [27], [28], [29],

[30], [31], [32], [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], has attracted the attention of the computational imaging community again recently due to recent advancements in deep generative modeling [44] and BNNs [17], [18], [20]. The state-of-the-art deep learning-based image reconstruction methods performing uncertainty characterization, e.g., [22], [24], [27], [28], [29], [30], [31], [32], [33], [34], [35], [37], [38], [39], [40], [41], [42], [43], can be divided into two groups: deep generative model-based reconstruction methods and BNN-based reconstruction methods.

Deep generative model-based reconstruction methods, e.g., [22], [24], [28], [29], [40], [42], seek to approximate the posterior distribution of the target image with the help of a generative model to characterize the inherent uncertainty in the reconstruction task, i.e., the uncertainty on the target image for a given measurement vector. For example, Adler and Oktom [22] approximate the posterior distribution of the target image given a measurement vector using a conditional Wasserstein generative adversarial network [45], [46]. Bohm et al. [24] use a variational autoencoder [47] to represent the prior distribution of the target image and perform variational inference to learn the true posterior distribution of the latent variable given a measurement vector. Sun and Bouman [40] utilize another popular generative model, a flow-based model [48], [49], to approximate the posterior distribution of the target image given a measurement vector and adjust the parameters of the flow-based model by minimizing the reverse Kullback-Leibler divergence [49] between the output distribution of the flow-based model and the posterior distribution. After training the generative model, the uncertainty on the target image for a given measurement vector can be quantified by calculating the sample variance of the samples generated from the approximation of the posterior distribution of the target image.

While deep generative model-based reconstruction methods aim to quantify the inherent uncertainty in the reconstruction task, the goal of Bayesian neural network-based reconstruction methods, e.g., [27], [30], [31], [33], [37], [38], [39], [41], [43], is to quantify either the uncertainty on the parameters of the statistical model or both the inherent uncertainty in the reconstruction task and the uncertainty on the parameters of the statistical model. To the best of the authors' knowledge, Schlemper et al. [37] presented the first two BNN-based image reconstruction methods for the MRI reconstruction problem, showing the potential of BNNs for uncertainty quantification for imaging problems. Subsequently, many BNN-based image reconstruction methods were developed for various problems such as the neuroimage enhancement [41], phase imaging [43], seismic imaging [39], computational optical form measurements [33], single-pixel imaging [38], imaging through scattering media [27], and more general image reconstruction problems [30], [31]. Table I shows the functional models of and the types of uncertainties quantified by BNN-based image reconstruction methods.

Table I highlights the main differences between the proposed framework and the aforementioned BNN-based image reconstruction methods. The main difference between the proposed framework and the methods presented in [27], [33], [37], [38], [39], [41], [43] is that the proposed framework utilizes the idea

TABLE I  
HIGH-LEVEL COMPARISON OF BAYESIAN NEURAL NETWORK-BASED IMAGE RECONSTRUCTION METHODS

Method	Functional Model	Quantified Uncertainties
Schlemper <i>et al.</i> [37]	U-Net [19]	Epistemic & Aleatoric
Schlemper <i>et al.</i> [37]	DCCNN [50]	Epistemic & Aleatoric
Tanno <i>et al.</i> [41]	ESPCN [51]	Epistemic & Aleatoric
Xue <i>et al.</i> [43]	U-Net [19]	Epistemic & Aleatoric
Siahkoohi <i>et al.</i> [39]	DIP [52]	Epistemic
Hoffmann <i>et al.</i> [33]	U-Net [19]	Epistemic
Shang <i>et al.</i> [38]	U-Net [19]	Epistemic & Aleatoric
Ekmekci and Cetin [30]	DRUNet [53]	Epistemic
Ekmekci and Cetin [31]*	Deep Unrolling	Epistemic
Cochrane <i>et al.</i> [27]	U-Net [19]	Epistemic
Proposed Framework	Deep Unrolling	Epistemic & Aleatoric

\*Preliminary version of this work

of deep unrolling to integrate the observation model into the reconstruction process. Incorporation of physics-based models through data-consistency layers provides some level of interpretability. The DCCNN [50] based method presented in [37] contains data-consistency layers; however, the data-consistency layer in [37] leverages the characteristic properties of the forward operator of the MRI observation model, making it highly specialized for MRI reconstruction. On the other hand, the proposed framework only requires the computation of the adjoint of the forward operator of the observation model, which is a considerably less restrictive requirement. If the forward operator deviates from a Fourier operator, the data consistency layer of the DCCNN-based method requires matrix inversion, which is not computationally feasible for large scale inverse problems. The difference between the proposed framework and the framework presented in [30] lies in the difference between end-to-end models and Plug-and-Play (PnP) methods [54]. While the BNN-based image reconstruction method presented in [30] is built upon the idea of Plug-and-Play (PnP) priors [54], which does not require end-to-end training, the proposed framework uses a deep unrolled network as its functional model, which is trained in an end-to-end manner.

We note that the preliminary version of this work appeared in [31] as a conference paper. The work presented in this manuscript extends the preliminary work in [31] in several significant ways. First, [31] involved the quantification of epistemic uncertainty only, whereas this article proposes both epistemic and aleatoric uncertainty quantification. Second, unlike [31], the unrolled neural network in the framework we propose here contains different CNN blocks at each iteration. We have experimentally observed that this change leads to a faster and more stable training stage. Finally, this manuscript contains an extensive set of experiments demonstrating the characteristics of epistemic and aleatoric uncertainties.

### III. PROPOSED FRAMEWORK

In this section, we present a learning-based image reconstruction framework that can incorporate the observation model into the reconstruction process and quantify epistemic and aleatoric uncertainties arising in imaging problems. We start by introducing the assumed observation model and presenting a

probabilistic formulation of deep unrolling methods along with a motivation for bringing in BNNs. This provides the basis for our BNN-based image reconstruction and uncertainty characterization approach, the components of which are described in the rest of this section.

#### A. Observation Model and the Inverse Problem

We consider the following observation model.

$$\mathbf{m} = \mathbf{As} + \mathbf{n}, \quad (1)$$

where  $\mathbf{m} \in \mathbb{F}^M$  is the measurement vector;  $\mathbf{A} \in \mathbb{F}^{M \times N}$  is the forward operator, which is the discrete approximation of the transformation applied by the imaging system;  $\mathbf{s} \in \mathbb{F}^N$  is the target image; and  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$  is additive white Gaussian noise, where  $\mathbb{F}$  stands for either  $\mathbb{R}$  or  $\mathbb{C}$ . In this section, without loss of generality, we only consider the case where  $\mathbb{F} = \mathbb{R}$  since generalizing the proposed framework to cover the case  $\mathbb{F} = \mathbb{C}$  is straightforward (see [31] for details).

For an underdetermined system ( $M < N$ ), the inverse problem, i.e., recovering the target image  $\mathbf{s}$  from the measurement vector  $\mathbf{m}$ , is an ill-posed problem. To narrow down the solution space, we can utilize any prior knowledge about the target image. One way to use such prior knowledge systematically is to treat the inverse problem as a maximum *a posteriori* (MAP) estimation problem, which is defined by

$$\hat{\mathbf{s}} = \underset{\mathbf{s} \in \mathbb{R}^N}{\operatorname{argmin}} \{ \|\mathbf{As} - \mathbf{m}\|_2^2 + \beta \psi(\mathbf{s}) \}, \quad (2)$$

where  $\hat{\mathbf{s}}$  is the MAP estimate of the target image, the term  $\|\mathbf{As} - \mathbf{m}\|_2^2$  is the data-fidelity term, the function  $\psi : \mathbb{R}^N \rightarrow \mathbb{R}$  is the regularizer that comes from the prior knowledge on the target image, and  $\beta > 0$  is the parameter controlling the balance between the data-fidelity term and the regularizer. After deciding on the form of the regularizer, e.g., total variation semi-norm or wavelet transform domain regularization, model-based reconstruction methods solve the problem in (2) iteratively by using an appropriate optimization algorithm, e.g., ADMM [6], HQS, or PGD [7].

#### B. Probabilistic Formulation of Deep Unrolling and BNNs

For the inverse problem, which is essentially a regression problem, suppose that the likelihood function  $p(\mathbf{s}|\mathbf{m}, \theta)$  has the following form.

$$p(\mathbf{s}|\mathbf{m}, \theta) = \mathcal{N}(\mathbf{s}|f_\theta(\mathbf{m}), \sigma^2 \mathbf{I}), \quad (3)$$

where  $f_\theta : \mathbb{R}^M \rightarrow \mathbb{R}^N$  is a deep unrolled network parametrized by the set of parameters  $\theta$ , and  $\sigma > 0$  is a fixed constant. Assuming that the training dataset  $\mathcal{D}$  contains i.i.d. pairs of measurement vectors and target images, we can compute a MAP estimate of the set of parameters by solving the following optimization problem.

$$\hat{\theta}_{\text{MAP}} = \underset{\theta}{\operatorname{argmin}} \left\{ \frac{1}{2\sigma^2} \sum_{i=1}^{N_D} \|\mathbf{s}^{[i]} - f_\theta(\mathbf{m}^{[i]})\|_2^2 - \log p(\theta) \right\} \quad (4)$$

where  $(\mathbf{m}^{[i]}, \mathbf{s}^{[i]})$  is the  $i^{\text{th}}$  example in the training dataset,  $N_{\mathcal{D}}$  is the number of examples in the training dataset, and the distribution  $p(\theta)$  is the prior distribution of the set of parameters. In the inference stage, for a given measurement vector  $\mathbf{m}_*$ , we can compute the distribution  $p(\mathbf{s}_* | \mathbf{m}_*, \hat{\theta}_{\text{MAP}})$  to make predictions about the target image.

This probabilistic formulation implicitly appears in the training and inference stages of state-of-the-art deep unrolling methods. For instance, if we choose the prior  $p(\theta)$  to be standard Gaussian distribution, then finding the MAP estimate of the set of parameters boils down to training the neural network  $f_{\theta}$  using the squared error loss with weight decay, which is a cost function frequently used by deep unrolling methods. In the inference stage, for a given measurement vector  $\mathbf{m}_*$ , outputting the mean of the distribution  $p(\mathbf{s}_* | \mathbf{m}_*, \hat{\theta}_{\text{MAP}})$  as the reconstructed image is equivalent to feeding the measurement vector  $\mathbf{m}_*$  into the trained neural network  $f_{\hat{\theta}_{\text{MAP}}}$ . Hence, training and inference procedures followed by many existing deep unrolling methods can be interpreted probabilistically using the formulation above.

Although such procedures are frequently used to train deep unrolling methods, there are two problems with this approach regarding the characterization of uncertainties. The first problem is that this formulation does not model the uncertainty on the target image for a given measurement vector, i.e., the inherent uncertainty on the reconstruction task, since it assumes that the covariance matrix of the likelihood function is a fixed model parameter. The second problem is that this formulation does not model the uncertainty on the set of parameters because it only uses a point estimate of the set of parameters by following MAP estimation principles.

BNNs [17], [18], [55] can solve these two problems. BNNs solve the first problem by defining a likelihood function that models the inherent uncertainty on the reconstruction task. In the case of a Gaussian likelihood, this can be accomplished by representing the covariance matrix of the likelihood function with a neural network. To solve the second problem, BNNs place a prior distribution on the set of parameters of the likelihood function and compute the posterior distribution of the parameters given a training dataset. Then, at the inference stage, BNNs compute the predictive distribution for a given measurement vector  $\mathbf{m}_*$  by computing the following integral:

$$p(\mathbf{s}_* | \mathbf{m}_*, \mathcal{D}) = \int p(\mathbf{s}_* | \mathbf{m}_*, \theta) p(\theta | \mathcal{D}) d\theta, \quad (5)$$

where the distribution  $p(\mathbf{s}_* | \mathbf{m}_*, \mathcal{D})$  is the predictive distribution, and the integral is taken over all possible values of  $\theta$ . The first term of the integrand, which is the likelihood function, incorporates the inherent uncertainty on the reconstruction task (i.e., aleatoric uncertainty), which is created by the ill-posedness of the inverse problem, into the predictive distribution. The second term of the integrand, on the other hand, which is the posterior distribution of the parameters, incorporates the uncertainty on the set of parameters (i.e., epistemic uncertainty), which is created by the lack of training examples in the training dataset around the test measurement vector, into the predictive distribution through an integral over the possible parameter values. Thanks to this conceptually simple probabilistic formulation,

we can utilize BNNs to quantify both epistemic and aleatoric uncertainties in computational imaging problems.

### C. Form of the Likelihood Function

Based on our observations presented in Section III-B, we define the form of likelihood function as follows.

$$p(\mathbf{s} | \mathbf{m}, \theta) = \mathcal{N}(\mathbf{s} | f_{\gamma}(\mathbf{m}), \text{diag}(\sigma_{\delta}^2(\mathbf{m}))), \quad (6)$$

where  $\theta = \gamma \cup \delta$ , and  $f_{\gamma} : \mathbb{R}^M \rightarrow \mathbb{R}^N$  and  $\sigma_{\delta}^2 : \mathbb{R}^M \rightarrow \mathbb{R}^N$  are two neural networks parametrized by sets of parameters  $\gamma$  and  $\delta$ , respectively. The neural network  $f_{\gamma}$  maps a given measurement vector to a point in the target image space, and the neural network  $\sigma_{\delta}^2$  aims to capture the inherent uncertainty on the target image for a given measurement vector. The form of the likelihood function implicitly assumes that  $\mathbf{s} = f_{\gamma}(\mathbf{m}) + \mathbf{n}_l$ , where  $\mathbf{n}_l \sim \mathcal{N}(\mathbf{0}, \text{diag}(\sigma_{\delta}^2(\mathbf{m})))$ , i.e., the aleatoric uncertainty is modeled as additive Gaussian noise. This assumption may be restrictive for severely ill-posed inverse problems since the distribution of the underlying image  $\mathbf{s}$  for a given measurement vector  $\mathbf{m}$  can be quite multimodal. To circumvent this issue, the likelihood function can be defined as a mixture of Gaussians distribution by following the idea of mixture density networks [56, Chapter 5.6]; however, this does not scale well to most imaging inverse problems since it requires training several neural networks (in the order of the number of mixture components). Another alternative is to introduce a latent variable for the deep neural network  $f$  to model complex uncertainty patterns on the target image, which is referred to as Bayesian neural networks with latent variables [57]. However, most deep unrolled networks are not latent variable models, and it is not straightforward to integrate latent variables into the deep unrolled architectures. Due to these reasons, we have decided to model the aleatoric uncertainty as additive Gaussian noise using the neural network  $\sigma_{\delta}^2$ .

To incorporate the observation model into the neural network  $f_{\gamma}$ , which maps a given measurement vector to a point on the target image space, we start constructing  $f_{\gamma}$  by first solving the optimization problem in (2) using the proximal gradient descent (PGD) method. The main advantage of using PGD over methods such as ADMM and HQS is that the data dependent update equation of PGD requires computing only the adjoint of the forward operator and does not involve any inversion step, which makes it suitable for large scale imaging problems with non-structured forward operators. Assuming that the regularizer  $\psi$  in (2) is a closed proper convex function, PGD yields the following iterative image reconstruction algorithm.

$$\begin{aligned} \mathbf{z}^{(k+1)} &= (\mathbf{I} - 2\alpha \mathbf{A}^{\top} \mathbf{A}) \mathbf{s}^{(k)} + 2\alpha \mathbf{A}^{\top} \mathbf{m} \\ \mathbf{s}^{(k+1)} &= \text{prox}_{\alpha \beta \psi} (\mathbf{z}^{(k+1)}) \end{aligned} \quad (7)$$

where  $\mathbf{z}^{(k+1)} \in \mathbb{R}^N$  is an intermediate vector of the algorithm at the  $(k+1)^{\text{st}}$  iteration,  $\mathbf{s}^{(k+1)} \in \mathbb{R}^N$  is the reconstructed image at the  $(k+1)^{\text{st}}$  iteration, the operator  $\text{prox} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is the proximal operator [7], and  $\alpha > 0$  is the step size. To learn the prior information implicitly from the training data, we replace the proximal operator in the second step with a neural network, which has been frequently done by deep unrolling methods such

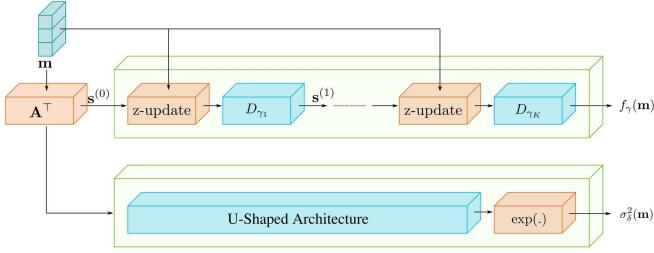


Fig. 1. The structure of the neural networks  $f_\gamma$  and  $\sigma_\delta^2$  at a high level. The neural network  $f_\gamma$  maps a measurement vector to a point in the target image space, and the neural network  $\sigma_\delta^2$  aims to capture the aleatoric uncertainty. These two neural networks completely specify the form of the Gaussian likelihood in (6).

as [16]. Then, the resulting update equations become

$$\begin{aligned} \mathbf{z}^{(k+1)} &= (\mathbf{I} - 2\alpha \mathbf{A}^\top \mathbf{A}) \mathbf{s}^{(k)} + 2\alpha \mathbf{A}^\top \mathbf{m} \\ \mathbf{s}^{(k+1)} &= D_{\gamma_{k+1}}(\mathbf{z}^{(k+1)}), \end{aligned} \quad (8)$$

where  $D_{\gamma_{k+1}} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is a residual neural network [58] parametrized by the set of parameters  $\gamma_{k+1}$ . For a fixed number of iterations  $K$ , the series of update equations in (8) correspond to a deep neural network  $f_\gamma$ , where  $\gamma = \bigcup_{k=1}^K \gamma_k$ . Fig. 1 displays the high-level summary of the neural network  $f_\gamma$ , and the details of the architecture are provided in the Supplementary Material.

To completely specify the form of the likelihood function given in (6), we have to specify the architecture of the neural network  $\sigma_\delta^2$  as well. The neural network we use for the neural network  $\sigma_\delta^2$  is a U-shaped neural network [19] followed by an element-wise exponentiation to ensure that the output contains positive entries. Fig. 1 depicts a high-level summary of the neural network  $\sigma_\delta^2$ , the details of which are given in the Supplementary Material. We must remark that we can also use a dual-head architecture to jointly represent the neural networks  $f_\gamma$  and  $\sigma_\delta^2$ . A brief discussion on the dual-head variant of the proposed framework is also provided in the Supplementary Material for interested readers.

#### D. Approximating the Posterior Distribution

To be able to compute the predictive distribution using (5), we have to compute the posterior distribution  $p(\theta|\mathcal{D})$ . However, exact computation of the posterior distribution is not tractable for deep neural networks because of the massive number of parameters and complex hierarchical structures. Thus, we either have to approximate the posterior distribution with a parametric distribution, or we have to generate samples from the posterior distribution to approximate the integral in (5). In our framework, we use a variational inference method called MC Dropout to approximate the posterior distribution with a parametric distribution. The advantages of using MC Dropout are that it is scalable for deep neural networks since it does not introduce additional parameters, variational inference and inference procedures can be straightforwardly implemented in deep learning frameworks because this only requires small changes on the training and testing procedures of the standard neural network

pipelines, and it has been shown that MC Dropout provides reliable uncertainty estimates for several problems such as camera relocalization [59], depth completion [18], [60] and semantic segmentation [18], [60].

For the sake of completeness, we state the assumptions of MC Dropout explicitly and discuss the variational inference and inference steps. For a more detailed discussion, the reader can refer to [18], [20], [61]. Suppose that the neural networks  $f_\gamma$  and  $\sigma_\delta^2$  contain  $C$  and  $E$  convolutional layers, respectively. Then, we can write the two sets  $\gamma$  and  $\delta$  as follows:

$$\gamma = \bigcup_{i=1}^C \{\mathbf{W}_i^f\} \quad \text{and} \quad \delta = \bigcup_{j=1}^E \{\mathbf{W}_j^\sigma\}, \quad (9)$$

where  $\mathbf{W}_i^f$  and  $\mathbf{W}_j^\sigma$  are the matrices whose rows contain the vectorized filter coefficients of the  $i^{\text{th}}$  and  $j^{\text{th}}$  convolutional layers of the neural networks  $f_\gamma$  and  $\sigma_\delta^2$ , respectively. The assumptions [18], [20], [61] on the parametric distribution  $q_\omega(\theta)$  that we use to approximate the true posterior distribution  $p(\theta|\mathcal{D})$  are as follows: i) For the parametric distribution, we assume that the layers of the neural networks  $f_\gamma$  and  $\sigma_\delta^2$  are independent, and layers within the neural networks are mutually independent, i.e.,

$$q_\omega(\theta) = \left( \prod_{i=1}^C q(\mathbf{W}_i^f) \right) \left( \prod_{j=1}^E q(\mathbf{W}_j^\sigma) \right); \quad (10)$$

ii) Filters of a convolutional layer are assumed to be mutually independent, more explicitly

$$q(\mathbf{W}_i^f) = \prod_{l=1}^{K_{i,f}^{[\text{out}]}} q\left(\left[\mathbf{W}_i^f\right]_{l,:}\right), \quad q(\mathbf{W}_j^\sigma) = \prod_{m=1}^{K_{j,\sigma}^{[\text{out}]}} q\left(\left[\mathbf{W}_j^\sigma\right]_{m,:}\right), \quad (11)$$

where  $K_{i,f}^{[\text{out}]}$  is the number of filters in the  $i^{\text{th}}$  convolutional layer of  $f_\gamma$ , and  $K_{j,\sigma}^{[\text{out}]}$  is the number of filters in the  $j^{\text{th}}$  convolutional layer of  $\sigma_\delta^2$ ; iii) The distribution of the filter coefficients of each filter is a mixture of Gaussians distribution defined by

$$\begin{aligned} q\left(\left[\mathbf{W}_i^f\right]_{l,:}\right) &= p(z_{i,l}^f = 1) \mathcal{N}\left(\left[\mathbf{W}_i^f\right]_{l,:} | \mathbf{a}_{i,l}^f, \epsilon^2 \mathbf{I}\right) \\ &\quad + p(z_{i,l}^f = 0) \mathcal{N}\left(\left[\mathbf{W}_i^f\right]_{l,:} | \mathbf{0}, \epsilon^2 \mathbf{I}\right), \\ q\left(\left[\mathbf{W}_j^\sigma\right]_{m,:}\right) &= p(z_{j,m}^\sigma = 1) \mathcal{N}\left(\left[\mathbf{W}_j^\sigma\right]_{m,:} | \mathbf{a}_{j,m}^\sigma, \epsilon^2 \mathbf{I}\right) \\ &\quad + p(z_{j,m}^\sigma = 0) \mathcal{N}\left(\left[\mathbf{W}_j^\sigma\right]_{m,:} | \mathbf{0}, \epsilon^2 \mathbf{I}\right), \end{aligned} \quad (12)$$

where the variables  $z_{i,l}^f$  and  $z_{j,m}^\sigma$  are the latent variables, and the scalars  $p_{i,l}^f \triangleq p(z_{i,l}^f = 1)$  and  $p_{j,m}^\sigma \triangleq p(z_{j,m}^\sigma = 1)$  are fixed constants. The scalar  $\epsilon$  is a very small fixed constant, and the sets  $\Delta_f \triangleq \{\mathbf{a}_{i,l}^f\}$  and  $\Delta_\sigma \triangleq \{\mathbf{a}_{j,m}^\sigma\}$  are the adjustable parameters of the parametric distribution. Previously we have denoted the set of adjustable parameters of the parametric distribution  $q_\omega(\theta)$  with  $\omega$ , so we can write the set  $\omega$  explicitly as  $\omega = \Delta_f \cup \Delta_\sigma$ .

Under these assumptions, we adjust the parameters of the parametric distribution by minimizing the Kullback-Leibler divergence between the parametric distribution and the true posterior distribution, i.e.,

$$\hat{\omega} = \operatorname{argmin}_{\omega} D_{\text{KL}}(q_{\omega}(\theta) \parallel p(\theta | \mathcal{D})). \quad (13)$$

Under certain approximations and mathematical manipulations (see the Supplementary Material or the Appendix of [20] for the details), the above optimization problem can be approximated with the following optimization problem.

$$\hat{\omega} \approx \operatorname{argmin}_{\omega} \{g(\omega) + h(\omega)\}, \quad (14)$$

where

$$\begin{aligned} g(\omega) &\triangleq \frac{1}{N_{\mathcal{D}}} \sum_{n=1}^{N_{\mathcal{D}}} \sum_{k=1}^N \left[ \frac{1}{2} \log[\sigma_{\tilde{\delta}^{(n)}}^2(\mathbf{m}^{[n]})]_k \right. \\ &\quad \left. + \frac{1}{2} \exp(-\log[\sigma_{\tilde{\delta}^{(n)}}^2(\mathbf{m}^{[n]})]_k) ([\mathbf{s}^{[n]}]_k - [f_{\tilde{\gamma}^{(n)}}(\mathbf{m}^{[n]})]_k)^2 \right], \\ h(\omega) &\triangleq \sum_{i=1}^C \sum_{l=1}^{K_{i,f}^{[\text{out}]} \atop K_{j,\sigma}^{[\text{out}]}} \frac{p_{i,l}^f}{2N_{\mathcal{D}}} \|\mathbf{a}_{i,l}^f\|_2^2 + \sum_{j=1}^E \sum_{m=1}^{K_{j,\sigma}^{[\text{out}]} \atop K_{j,m}^{[\text{out}]}} \frac{p_{j,m}^{\sigma}}{2N_{\mathcal{D}}} \|\mathbf{a}_{j,m}^{\sigma}\|_2^2, \end{aligned} \quad (15)$$

and  $\tilde{\theta}^{(n)} = \tilde{\delta}^{(n)} \cup \tilde{\gamma}^{(n)}$  is the  $n^{\text{th}}$  sample generated from the parametric distribution  $q_{\omega}(\theta)$ .

After approximating the true posterior distribution  $p(\theta | \mathcal{D})$  with the parametric distribution  $q_{\hat{\omega}}(\theta)$ , we approximate the integral in (5) using Monte Carlo integration with  $T$  samples as follows.

$$p(\mathbf{s}_* | \mathbf{m}_*, \mathcal{D}) \approx \frac{1}{T} \sum_{t=1}^T \mathcal{N} \left( \mathbf{s}_* | f_{\tilde{\gamma}^{(t)}}(\mathbf{m}_*), \operatorname{diag} \left( \sigma_{\tilde{\delta}^{(t)}}^2(\mathbf{m}_*) \right) \right) \quad (16)$$

where  $\hat{\theta}^{(t)} = \hat{\delta}^{(t)} \cup \hat{\gamma}^{(t)}$  is the  $t^{\text{th}}$  sample from the parametric distribution  $q_{\hat{\omega}}(\theta)$ . The approximation of the predictive distribution is a mixture of  $T$  Gaussians with uniform weights; therefore, we can compute its mean vector and element-wise variance analytically as follows.

$$\mathbb{E}[\mathbf{s}_* | \mathbf{m}_*, \mathcal{D}] \approx \frac{1}{T} \sum_{t=1}^T f_{\tilde{\gamma}^{(t)}}(\mathbf{m}_*), \quad (17)$$

$$\begin{aligned} \operatorname{Var}[[\mathbf{s}_*]_k | \mathbf{m}_*, \mathcal{D}] &\approx \underbrace{\frac{1}{T} \sum_{t=1}^T \left[ \sigma_{\tilde{\delta}^{(t)}}^2(\mathbf{m}_*) \right]_k}_{\text{Aleatoric variance}} \\ &\quad + \underbrace{\frac{1}{T} \sum_{t=1}^T \left[ f_{\tilde{\gamma}^{(t)}}(\mathbf{m}_*) \right]_k^2 - \left( \frac{1}{T} \sum_{t=1}^T \left[ f_{\tilde{\gamma}^{(t)}}(\mathbf{m}_*) \right]_k \right)^2}_{\text{Epistemic variance}}, \end{aligned} \quad (18)$$

where  $\hat{\theta}^{(t)} = \hat{\delta}^{(t)} \cup \hat{\gamma}^{(t)}$  is the  $t^{\text{th}}$  sample from the optimized parametric distribution  $q_{\hat{\omega}}(\theta)$ . The first term of the predictive variance, which we refer to as the aleatoric variance, reflects the aleatoric uncertainty in the reconstruction task, and the

remaining residual sum, which we refer to as the epistemic variance, represents the epistemic uncertainty.

At this point, we have to be aware that we have treated the parameters of the neural networks  $f_{\gamma}$  and  $\sigma_{\delta}^2$  as random variables and have to generate samples from the parametric distribution to solve the optimization problem in (14) and to obtain the predictive mean and variance given by (17) and (18). Because we have assumed that filters of convolutional layers are mutually independent, one naive way to generate a sample from the parametric distribution is to generate samples from the distributions in (12). Sampling from those distributions is equivalent to sampling from a mixture of Gaussians distribution with two components, so, first, we need to sample a Bernoulli random variable, and based on that sample, we generate a sample from one of the two multivariate Gaussian distributions. Because the scalar  $\epsilon$  is assumed to be a very small non-zero constant, generating a sample from the multivariate Gaussian distributions in (12) can be approximated by directly reporting the mean. Thus, the whole process of generating a sample from the parametric distribution  $q_{\omega}(\theta)$  boils down to generating samples from Bernoulli random variables and multiplying them with the adjustable parameters of the parametric distribution. Hence we can write

$$\begin{aligned} \tilde{\gamma}^{(n)} &\approx \left\{ \tilde{z}_{i,l}^{(n)} \mathbf{a}_{i,l}^f \mid \text{sample } \tilde{z}_{i,l}^{(n)} \sim \text{Bernoulli} \left( p_{i,l}^f \right) \right\}, \\ \tilde{\delta}^{(n)} &\approx \left\{ \tilde{z}_{j,m}^{(n)} \mathbf{a}_{j,m}^{\sigma} \mid \text{sample } \tilde{z}_{j,m}^{(n)} \sim \text{Bernoulli} \left( p_{j,m}^{\sigma} \right) \right\}, \\ \tilde{\theta}^{(n)} &= \tilde{\delta}^{(n)} \cup \tilde{\gamma}^{(n)}. \end{aligned} \quad (19)$$

An interesting observation is that the sampling operation described above resembles the dropout [62] operation. Hence, solving the optimization problem in (14) boils down to training two neural networks  $f$  and  $\sigma^2$  only once using the loss function defined by

$$\begin{aligned} L(\omega) &= \frac{1}{N_{\mathcal{B}}} \sum_{n=1}^{N_{\mathcal{B}}} \sum_{k=1}^N \left[ \frac{1}{2} \log[\bar{\sigma}_{\Delta_{\sigma}}^2(\mathbf{m}^{[n]})]_k \right. \\ &\quad \left. + \frac{1}{2} \exp \left( -\log \left[ \bar{\sigma}_{\Delta_{\sigma}}^2(\mathbf{m}^{[n]}) \right]_k \right) \right. \\ &\quad \left. \times \left( [\mathbf{s}^{[n]}]_k - [\bar{f}_{\Delta_f}(\mathbf{m}^{[n]})]_k \right)^2 \right], \end{aligned} \quad (20)$$

where  $\mathcal{B}$  is a mini-batch from the training dataset  $\mathcal{D}$ , and  $N_{\mathcal{B}}$  is the size of the mini-batch, with weight decay parameters  $p_{i,l}^f / (2N_{\mathcal{D}})$  and  $p_{j,m}^{\sigma} / (2N_{\mathcal{D}})$  and with dropout rates  $1 - p_{i,l}^f$  and  $1 - p_{j,m}^{\sigma}$ . After the training stage, the resulting weights of the dropout-added neural networks will be the optimal parameters  $\hat{\omega}$  of the parametric distribution  $q_{\hat{\omega}}(\theta)$ , and generating a sample from the parametric distribution  $q_{\hat{\omega}}(\theta)$  simply boils down to applying dropout to the weights of the dropout-added neural networks. Furthermore, calculating the approximation of the predictive distribution using (16) boils down to feeding the test measurement vector to the trained dropout-added neural networks  $\bar{f}_{\hat{\Delta}_f}$  and  $\bar{\sigma}_{\hat{\Delta}_{\sigma}}$  exactly  $T$  times while the dropout is on. To obtain a reconstruction, we can either generate samples from the approximation of the predictive distribution or compute its mean

using (17). To obtain the epistemic and aleatoric uncertainty maps, we use the expression in (18).

We must remark that we have not included the step size parameter  $\alpha$  used in the deep unrolled network  $f$  and the parameters of the batch normalization layers of the neural network  $\sigma^2$ , in the sets  $\gamma$  and  $\delta$ . In other words, we have not taken those parameters into account in the Bayesian formulation. If we want to include the parameters of a batch normalization layer in the formulation, we have to randomly set the parameters of the batch normalization layer to zero, which is not a common practice in deep learning literature. Thus, we decide to treat the parameters of the batch normalization layers as trainable deterministic parameters. Similarly, randomly setting the step size parameter to zero violates the positivity requirement of the step size. We could choose other, preferably more expressive, parametric distributions on the step size parameter to perform variational inference or use deep ensembling [63] to capture the uncertainty on the step size parameter. However, we have experimentally observed that treating the step size parameter and the parameters of the batch normalization layers as trainable deterministic parameters is enough to obtain meaningful uncertainty estimates without introducing more complexity to the model.

#### IV. EXPERIMENTS AND RESULTS

In this section, we present experimental results demonstrating the behavior of our proposed approach. Although the proposed framework can be applied to any inverse problem that can be cast as the optimization problem in (2), we evaluate the proposed framework on basic MRI and CT reconstruction problems as representative applications. We investigate the behavior of epistemic and aleatoric uncertainties under various experimental conditions and show that the epistemic and aleatoric uncertainty information provided by the proposed framework is consistent with the definitions of those uncertainties. We then investigate the calibration properties of the proposed method by generating calibration plots for the MRI and CT reconstruction problems. Finally, we compare the image reconstruction performance of the proposed framework with other image reconstruction methods to demonstrate the image reconstruction capability of the proposed framework. Supplementary Material also contains a toy problem illustrating the concepts of epistemic and aleatoric uncertainties.

##### A. Experimental Setup

**Datasets:** For the MRI reconstruction problem, we extracted 530  $256 \times 256$  target images from the IXI Dataset [64]. Each target image was normalized between 0 and 1. We split the 530 target images into training, validation, and test datasets containing 500, 15, and 15 target images, respectively. The training, validation, and test datasets were constructed such that they contain target images collected from different subjects. The measurement vectors, i.e., k-space measurements, were generated by computing the subsampled Fourier transform of the target images. For the CT reconstruction problem, we extracted 530  $512 \times 512$  target images from the LUNA Dataset [65].

Each image was resized to  $256 \times 256$  pixels and normalized between 0 and 1. The training dataset was created by using 500 target images, and the remaining 30 images were split into two sets to generate validation and test datasets, each containing 15 target images. The training, validation, and test datasets were constructed such that they contain target images collected from different subjects. The measurement vectors, i.e., sinogram data, were generated by computing the sparse Radon transform of the target images. Finally, we added white Gaussian noise to the measurement vectors to obtain the final measurement vectors we used in our experiments, where the SNR of the measurement vector is defined as follows:

$$\text{SNR}(\mathbf{m}_{\text{noiseless}} + \mathbf{n}, \mathbf{m}_{\text{noiseless}}) = 20 \log_{10} \left( \frac{\|\mathbf{m}_{\text{noiseless}}\|_2}{\|\mathbf{n}\|_2} \right). \quad (21)$$

**Training and Inference Procedures:** Training of the proposed framework refers to solving the optimization problem in (14), where the first term of the objective function is replaced with its mini-batch approximation [20]. We obtained the neural network  $\bar{f}$  by fixing the number of iterations  $K$  of the PGD to be 5 and taking the starting point  $\mathbf{s}^{(0)}$  to be the result of zero-filling and filtered backprojection for MRI and CT reconstruction problems, respectively. Each residual block of the neural network  $\bar{f}$  contains 5 convolutional layers, and each convolutional layer is followed by a dropout layer and the leaky ReLU activation function. We used the U-Net architecture for the neural network  $\bar{\sigma}^2$ , where each convolutional layer is followed by a dropout layer. For the MRI reconstruction problem, the batch size for the training was set to 4, and the learning rate was fixed to  $1 \times 10^{-4}$ . For the CT reconstruction problem, we used a batch size of 2 for the training and set the learning rate to  $1 \times 10^{-5}$ . The initial step size  $\alpha$  of the PGD algorithm was set to 1.0 for the MRI experiments and  $1 \times 10^{-4}$  for the CT experiments. The dropout rate of the dropout layers of the neural networks  $\bar{f}$  and  $\bar{\sigma}^2$  was set to 0.1, and the neural networks  $\bar{f}$  and  $\bar{\sigma}^2$  were trained for 100 epochs. At the inference stage, a given measurement vector was fed to the neural networks  $\bar{f}$  and  $\bar{\sigma}^2$   $T = 100$  times while the dropout was still activated. The reconstructed image was then obtained by following the approximation in (17). The epistemic and aleatoric uncertainty maps were obtained by calculating three times of the epistemic and aleatoric standard deviations given by (18).

##### B. Epistemic Uncertainty

In theory, epistemic uncertainty is the uncertainty created by the lack of training data around test data and can be explained away by making appropriate changes on the training data. In this subsection, we investigate the characteristics of epistemic uncertainty information provided by the proposed framework and show that the behavior and results of our approach are consistent with the theoretical characteristics of epistemic uncertainty.

To show that the proposed framework outputs epistemic uncertainty information that reflects the uncertainty caused by the lack of training data that can explain the test sample well, we consider two scenarios. In the first scenario, we assess the impact of the size of the training dataset on the inferred epistemic

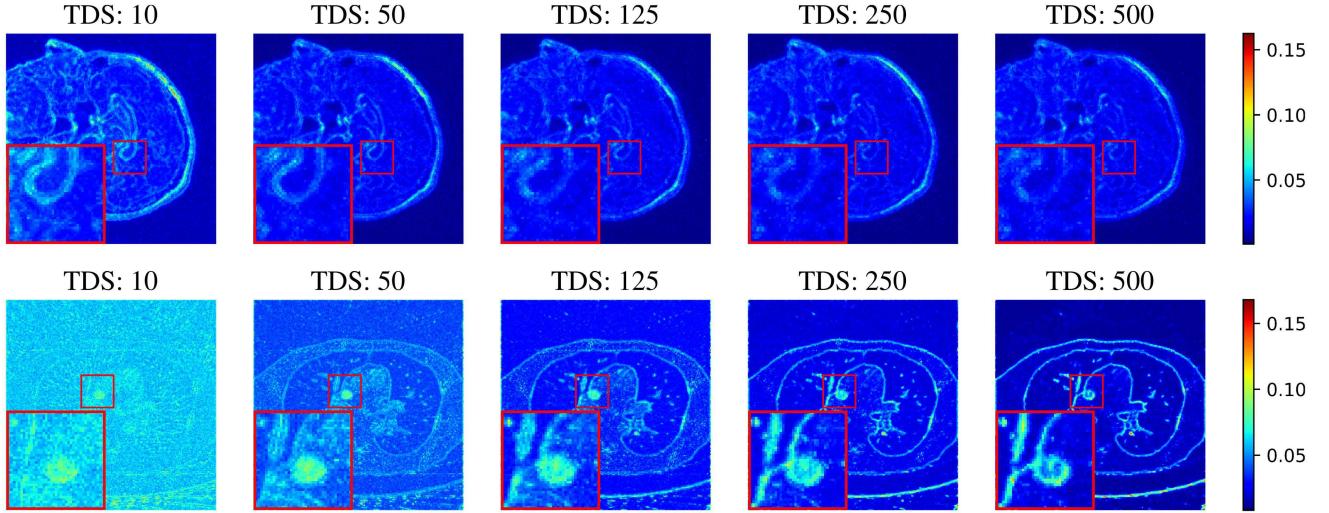


Fig. 2. Epistemic uncertainty maps on an MRI (top) and a CT (bottom) test sample as a function of the training dataset size (TDS). As we increase the number of examples in the training dataset, the overall epistemic uncertainty decreases. For the MRI experiments, the percentage of observed k-space coefficients is 20%, and SNR is 70 dB. For the CT experiments, number of views is 36, and SNR is 70 dB.

uncertainty. A good uncertainty characterization method should yield larger epistemic uncertainties for smaller training datasets, as it is less probable for such data to represent a random test sample well. In our experiments, we generated five subsets of the MRI training dataset containing 10, 50, 125, 250, and 500 examples, and trained five instances of the proposed framework using these subsets as training datasets. Then, for a given test measurement, we obtained the epistemic uncertainty maps using the five trained instances of the proposed framework. We repeated the same procedure for the CT reconstruction problem. The resulting epistemic uncertainty maps are illustrated in Fig. 2. For both MRI and CT reconstruction problems, epistemic uncertainty achieves its highest value for the case where we use only 10 training examples. Then, as we add more examples to the training dataset, epistemic uncertainty on the same test image decreases. To confirm these visual results quantitatively, we calculated the average epistemic uncertainty per pixel taken over the test dataset as a function of the size of the training dataset. Fig. 5 shows the quantitative results for both MRI and CT reconstruction problems. From this figure, we observe that an increase in the number of training examples leads to a decrease in the overall epistemic uncertainty, which is consistent with the visual results presented in Fig. 2. For a discussion about the effect of the size of the training dataset on the aleatoric uncertainty, please refer to the Supplementary Material.

For the second scenario, we insert an artificial feature that is not well-represented by the training dataset to a test target image. Then, we vary the intensity of the inserted abnormal feature to modify the degree of deviation of the test data from the training data. A good uncertainty characterization method would result in larger epistemic uncertainty as the test sample deviates more from the training data. In our experiments, we first trained the proposed framework on the MRI training dataset. Next, we picked a target image from the test dataset and inserted a  $25 \times 25$  square with the intensity value of 1.0 to the test target

image. Then, we obtained the epistemic uncertainty map. We repeated the same procedure for different values of the intensity of the inserted abnormal feature and for the CT reconstruction problem. Fig. 3 shows the epistemic uncertainty maps obtained by the proposed framework for different intensity values of the inserted abnormal feature for both MRI and CT reconstruction problems. We observe that the epistemic uncertainty in the abnormal region decreases as the intensity of the inserted square gets close to a value that makes the inserted square visually similar to the target images in the training dataset. Thus, our experiment shows that the epistemic uncertainty map obtained by the proposed framework shows high epistemic uncertainty for test data that are not well-represented by the training data, confirming that for this experiment, the proposed framework successfully captures the uncertainty caused by the lack of training data around the test data.

Now, we demonstrate that the epistemic uncertainty provided by the proposed framework possesses the reducibility property. For the first scenario, we have already shown in Figs. 2 and 5 that we can reduce the epistemic uncertainty by collecting more training data having similar characteristics to the test data. For the second scenario, if the proposed framework is capable of capturing the epistemic uncertainty well, we expect to see that using training examples containing features similar to the abnormal feature encountered in the test data would result in reduced epistemic uncertainty. To this end, we added  $25 \times 25$  white squares on the training target images and trained the proposed framework with such training data containing the abnormal features. We repeated the same procedure for the CT reconstruction problem. Fig. 4 shows the resulting epistemic uncertainty maps obtained by the proposed framework for both CT and MRI reconstruction problems. We observe that the epistemic uncertainty around the white square is decreased significantly after we added target images containing white squares into the training dataset, confirming that the epistemic uncertainty

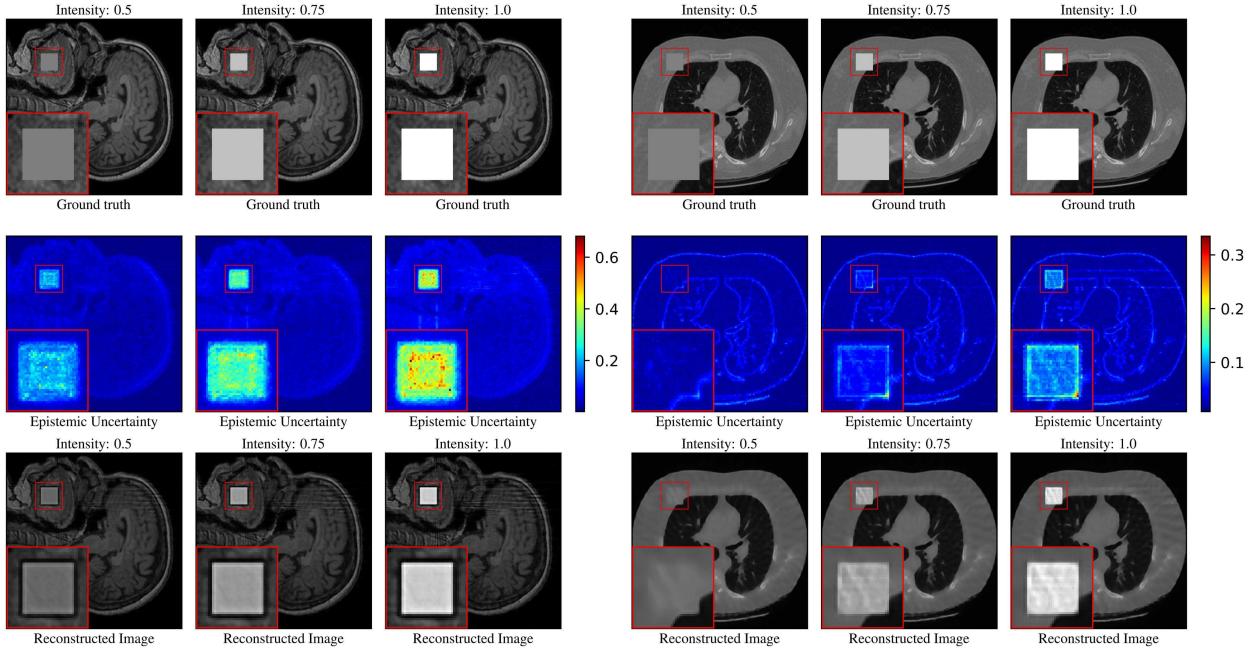


Fig. 3. Epistemic uncertainty as a function of the intensity of the abnormal feature. The first row contains the ground truth test images, i.e., the target test images. The second row contains the corresponding epistemic uncertainty maps obtained by the proposed framework, and the third row contains the corresponding reconstructed images. As the inserted square deviates more from the pattern of intensities in the test image (which would be well-represented by the training data), the inferred epistemic uncertainty in the abnormal region increases. For the MRI experiments, the percentage of observed k-space coefficients is 20%, and SNR is 70 dB. For the CT experiments, number of views is 36, and SNR is 70 dB.

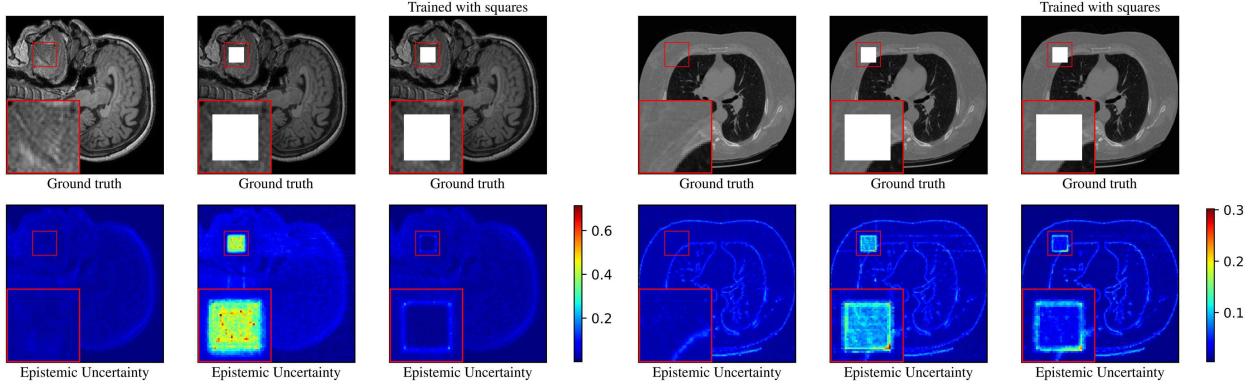


Fig. 4. Effect of the structure of the training dataset on epistemic uncertainty maps. The first row contains the ground truth test images, i.e., the target test images, and the second row contains the corresponding epistemic uncertainty maps. The images on the first and fourth columns show the performance of the proposed framework on normal data (i.e., no abnormal features in training and test data). The images on the second and fifth columns show the performance of the proposed framework on a case where an abnormal feature exists in the test data. The images on the third and sixth columns show the performance of the proposed framework with abnormal features present in both training and test data. For the MRI experiments, the percentage of observed k-space coefficients is 20%, and SNR is 70 dB. For the CT experiments, number of views is 36, and SNR is 70 dB.

provided by the proposed framework can be explained away with additional training data that can represent the test data well. These results qualitatively show that the epistemic uncertainty estimates provided by the proposed method display the reducibility property of the epistemic uncertainty.

### C. Aleatoric Uncertainty

We now focus on aleatoric uncertainty characterization using our proposed framework. The experiments we present here demonstrate successful aleatoric uncertainty characterization,

in particular, the aleatoric uncertainty captured by the proposed framework is high for the regions where the reconstruction is challenging due to the ill-posed nature of the inverse problem. Furthermore, we show that the overall aleatoric uncertainty provided by the proposed framework is an indication of how challenging the inverse problem is. For this analysis, we trained the proposed framework for various configurations of the imaging setups. We considered different percentages of observed k-space coefficients and SNR values for the MRI reconstruction problem and different number of views and SNR values for the CT reconstruction problem. Fig. 6 shows the starting points

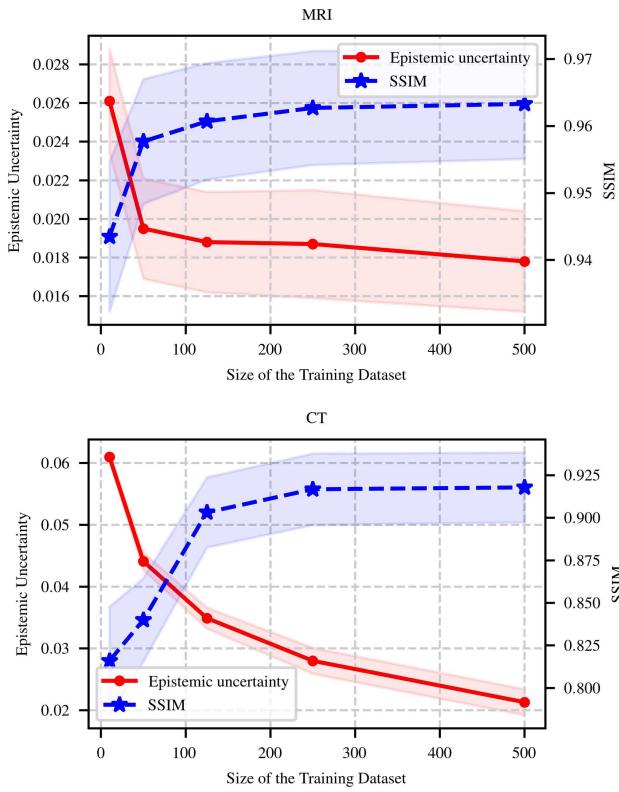


Fig. 5. Mean and standard deviation of epistemic uncertainty as a function of training dataset size. The mean and standard deviation are calculated using all pixels in the test dataset. For the MRI experiments, the percentage of observed k-space coefficients is 20%, and SNR is 70 dB. For the CT experiments, number of views is 36, and SNR is 70 dB. Mean SSIM values along with the standard deviations for the corresponding reconstructions are provided for reference.

of the proposed framework, i.e., the results of zero-filling and filtered backprojection, and the aleatoric uncertainty maps for different test measurement vectors generated from the two test target images using different configurations of the MR and CT imaging setups.

For both MRI and CT reconstruction problems, we observe that the aleatoric uncertainty is high for the regions where the reconstruction is challenging for the unrolled network, such as the small localized structures and thin edges on the target images. On the other hand, we observe that the aleatoric uncertainty is low around the regions where the corruption is negligible or can be recovered using the spatial information, such as the smooth regions in the target images. This behavior can be understood analytically with a careful inspection of the objective function of the optimization problem in (14). To minimize the objective function of the optimization problem given in (14), the optimization algorithm needs to minimize the term  $\log[\bar{\sigma}^2(\mathbf{m}^{[n]})]_k$  for the  $k^{\text{th}}$  pixel. However, the term  $\exp(-\log[\bar{\sigma}^2(\mathbf{m}^{[n]})]_k)$  would increase exponentially if the aforementioned term is minimized. Hence, the value of the term  $\log[\bar{\sigma}^2(\mathbf{m}^{[n]})]_k$  could be made small by the optimization algorithm if the squared error between the output of the neural network  $\bar{f}$  and the target image, i.e.,  $([\mathbf{s}^{[n]}]_k - [\bar{f}(\mathbf{m}^{[n]})]_k)^2$ , is small. In other words, solving optimization problem in (14), which corresponds to performing variational inference using MC Dropout, explicitly forces the

dropout added neural network  $\bar{\sigma}^2$  to output lower values where the reconstruction is relatively easy. On the other hand, for the regions where the reconstruction is challenging, i.e., for the regions where the squared error between the output of the neural network  $\bar{f}$  and the target image is high, solving this optimization problem forces the neural network  $\bar{\sigma}^2$  to output high values. Moreover, we observe that the overall aleatoric uncertainty levels show a decrease as SNR decreases for a fixed percentage of the observed k-space coefficients/number of views. Similarly, for a fixed value of the SNR, we observe a decrease in the overall aleatoric uncertainty levels as the percentage of the observed k-space coefficients/number of views increases. Fig. 7 shows the average aleatoric uncertainty over all pixels in the test dataset for different configurations of the imaging setups. From this figure, we observe that the overall aleatoric uncertainty increases when the SNR decreases for a fixed percentage of the observed k-space coefficients/number of views or when the percentage of the observed k-space coefficients/number of views decreases for a fixed value of the SNR. Hence, the quantitative results shown in Fig. 7 confirm our visual observations about the overall aleatoric uncertainty. This result can be also understood by analyzing the objective function of the optimization problem in (14). Because the neural network  $\bar{f}$  does not have an infinite learning capability in practice, we expect that the squared error between the output of the trained neural network  $\bar{f}$  and the target image will increase as the reconstruction problem gets more challenging, leading to higher overall aleatoric uncertainty levels for the relatively more challenging image reconstruction problems.

#### D. Calibration Plots of the Proposed Method

So far, we have observed that epistemic and aleatoric uncertainty maps convey useful information about the confidence of the reconstruction method and the imaging problem; however, we need to perform a more quantitative analysis to evaluate the probabilistic predictions of the proposed framework more reliably. One way of assessing the accuracy of the probabilistic predictions is to look at the calibration and the sharpness properties of the proposed model. In this subsection, we present calibration plots of the proposed method for the MRI and CT reconstruction problems. Furthermore, we briefly touch on the recalibration of the proposed method to achieve more calibrated probabilistic predictions. Since the sharpness metric is useful to compare two probabilistic models, we have included the discussion of the sharpness of the proposed method in the supplementary material, where we compare the probabilistic predictions of the uncertainty-quantifying PnP method presented in [30] with the proposed method.

A calibration plot is a diagnostic tool that allows visually inspecting the calibration properties of a probabilistic model to understand whether the model is providing underconfident or overconfident predictions. In this section, to obtain calibration plots, we approximate the predictive distribution of the proposed method, which is a mixture of Gaussians distribution with  $T$  mixture components, with a multivariate Gaussian distribution. More specifically, we approximate the predictive distribution of

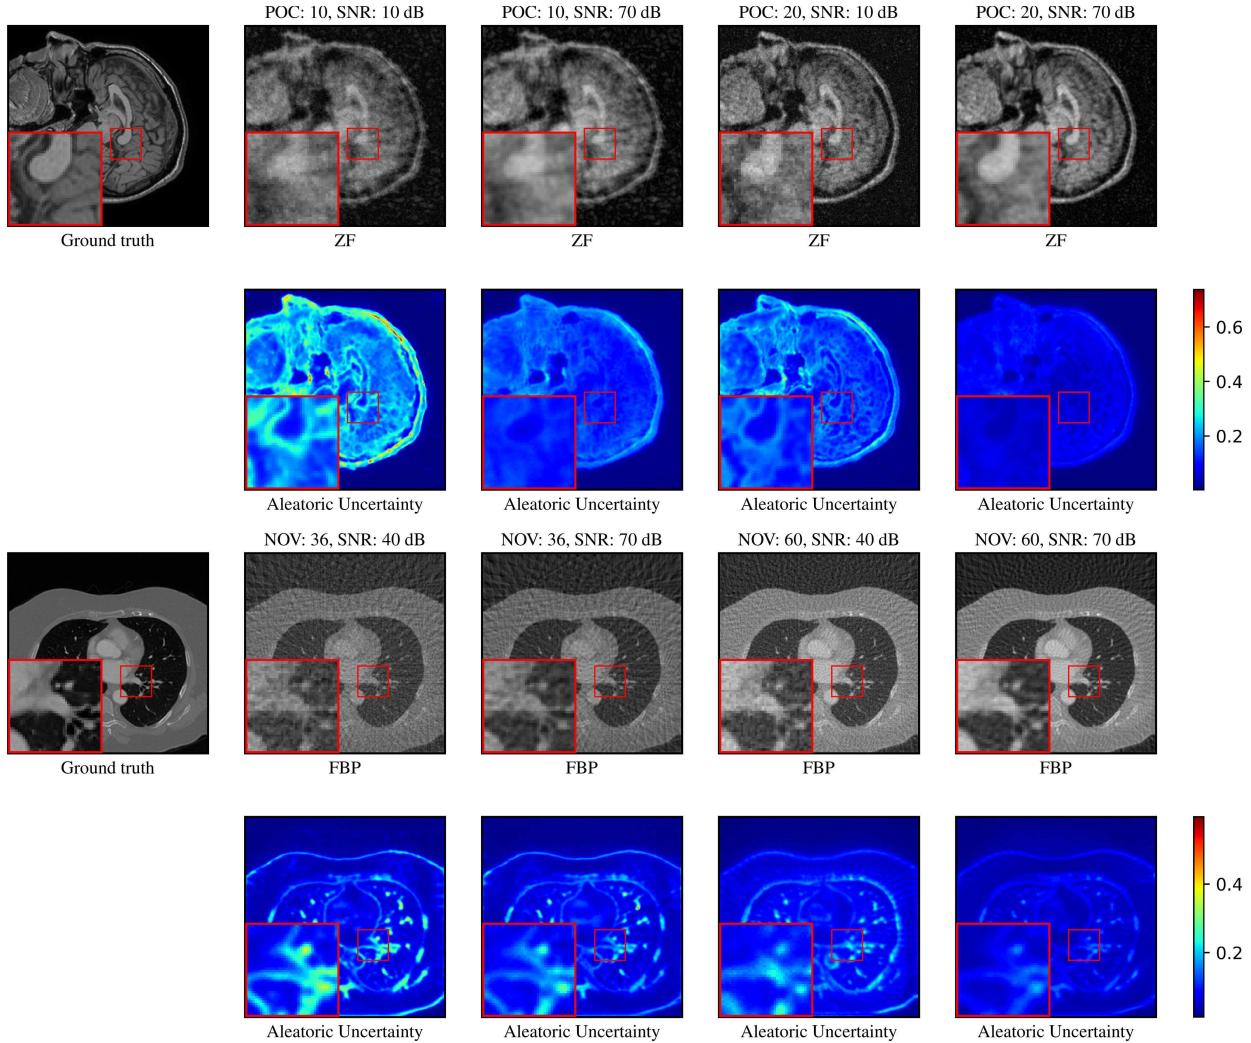


Fig. 6. Effect of the configuration of the imaging setup on aleatoric uncertainty. The first and the third rows contain the ground truth test images, i.e., target test images, as well as the starting points obtained by applying zero-filling (ZF) or filtered backprojection (FBP) to observations. The second and fourth rows contain the corresponding aleatoric uncertainty maps obtained by the proposed framework for different percentages of observed k-space coefficients (POC), numbers of views (NOV), and signal-to-noise ratios (SNR). Regions where the reconstruction from the starting point is challenging are the regions for which the aleatoric uncertainty is high. Moreover, the overall aleatoric uncertainty increases as the reconstruction problem gets more challenging in terms of data quality and quantity limitations.

each pixel with a Gaussian distribution as follows:

$$\begin{aligned} p([\mathbf{s}_*]_k | \mathbf{m}_*, \mathcal{D}) \\ \approx \mathcal{N}([\mathbf{s}_*]_k | [\mathbb{E}[\mathbf{s}_* | \mathbf{m}_*, \mathcal{D}]]_k, \text{Var}[[\mathbf{s}_*]_k | \mathbf{m}_*, \mathcal{D}]), \end{aligned} \quad (22)$$

where the mean and the variance of the distribution are defined in (17) and (18), respectively. Using this approximation, we generated calibration plots on the test datasets using Uncertainty Toolbox [66] for different configurations of the MRI and CT observation models. Figs. 8 and 9 show those calibration plots for the MRI and CT experiments, respectively. The red curves in the plots represent the calibration curves of the proposed method for different configurations of the imaging setups. The dashed green line, on the other hand, represents the *ideal* calibration curve.

Figs. 8 and 9 show that the proposed model may provide slightly underconfident predictions. The main reason behind this bias, which is sometimes referred to as the model bias, is the

assumptions we have made about the form of the likelihood function, the prior distribution of the parameters of the likelihood function, and the choice of the parametric distribution that we use to approximate the true posterior distribution of the parameters. Luckily, we can easily *recalibrate* the proposed method by following the recalibration method introduced by Kuleshov et al. [21]. In our experiments, we used the validation dataset as the calibration dataset to recalibrate the proposed method. Calibration curves of the calibrated proposed method are represented with the blue color in Figs. 8 and 9. After recalibration, we observe that the calibrated proposed model is capable of outputting more calibrated probabilistic predictions.

#### E. Reconstruction Performance

In this subsection, we demonstrate the reconstruction performance of the proposed framework. We compare the proposed

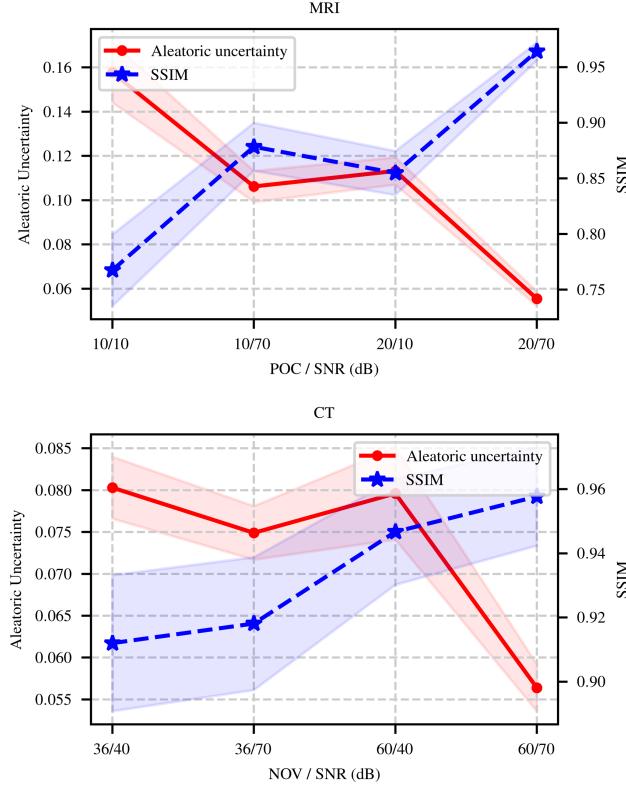


Fig. 7. Mean and standard deviation of aleatoric uncertainty for different configurations of the imaging setups. For MRI experiments, they are calculated for different percentages of observed k-space coefficients (POC) and signal-to-noise ratios (SNR). For CT experiments, they are calculated for different numbers of views (NOV) and signal-to-noise ratios (SNR). The averages and standard deviations are calculated using all pixels in the test dataset. Mean SSIM values along with the standard deviations for the corresponding reconstructions are provided for reference.

framework with six methods: 1) zero-filling (ZF) / filtered backprojection (FBP), 2) total variation reconstruction (TV), 3) PGD-based deep unrolling method (PUM), 4) PGD-based deep unrolling method without batch normalization (PUMw/oBN), 5) proposed only epistemic model (POEM), and 6) proposed only aleatoric model (POAM).

The methods ZF/FBP, and TV are the baseline reconstruction methods that we use to demonstrate how challenging the reconstruction problem is. PUM is a deep unrolling method using PGD. Each residual block of PUM consists of a series of convolutional layers, batch normalization layers, and an activation function. PUMw/oBN is the same model as the PUM, except that there are no batch normalization layers in residual blocks. POEM is the variant of the proposed framework that assumes that the covariance matrix of the likelihood function in (6) is a fixed model parameter. POEM is also the probabilistic model that was used in the experiments of the preliminary version of this article [31]. As its name implies, POEM quantifies only the epistemic uncertainty, not the aleatoric uncertainty. POAM is also a variant of the proposed framework where a maximum likelihood estimate of the parameters of the likelihood function in (6) is used. POAM is capable of quantifying the aleatoric uncertainty, but not the epistemic uncertainty since it only uses the MAP estimate of the parameters. Implementation details of these methods are provided in the Supplementary Material.

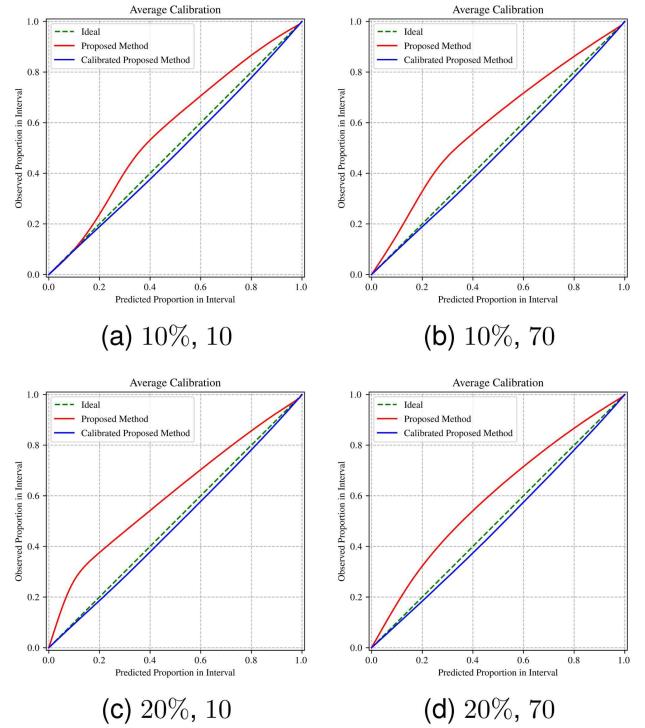


Fig. 8. Calibration plots of the proposed method for different configurations of the MRI setup. Subcaptions specify the percentage of observed k-space coefficients and the SNR (dB).

Table II shows the performance of the seven methods for CT and MRI reconstruction problems under different configurations of the imaging setups. Among these image reconstruction methods, FBP and ZF achieve the worst reconstruction performance among the seven reconstruction methods. The TV method improves upon FBP and ZF by promoting a piecewise-constant reconstruction. The deep unrolling method PUM surpasses the TV method by implicitly learning the prior using the training dataset. The deep unrolling method PUM was trained using a small mini-batch size since it requires storing the intermediate variables having the same spatial dimensions as the target image in the memory to carry out the backpropagation. We empirically observed that the removal of the batch normalization layers from the unrolled network leads to an increase in the reconstruction performance. Specifically, we observe that the PUMw/oBN outperforms PUM in all the experiments. This empirical observation is mathematically justified in [67] where Yong et al. showed that batch normalization introduces a high level of noise for small mini-batch sizes, making the training difficult. This observation is the main reason why the unrolled network  $f$  in the proposed framework does not contain any batch normalization layers. On the other hand, we experimentally observed that the addition of the batch normalization layers into the neural network  $\sigma^2$  is necessary to have a stable training stage. Comparing POAM with PUMw/oBN, POAM shows an average SSIM decrease of 0.022 for the MRI reconstruction problem and 0.002 for the CT reconstruction problem. On the other hand, when compared to the state-of-the-art deep unrolling method PUM, POAM achieves average SSIM gains of 0.031 and 0.011 for the MRI

TABLE II  
COMPARISON OF AVERAGE SSIM FOR DIFFERENT IMAGE RECONSTRUCTION METHODS

	POC	NOV	SNR	ZF	FBP	TV	PUM	PUMw/oBN	POAM	POEM	Proposed
MRI	10	-	10	0.4910	-	0.7033	0.7979	0.8144	0.7773	0.7727	0.7674
	10	-	70	0.5448	-	0.7261	0.8032	0.9227	0.8996	0.8638	0.8784
	20	-	10	0.5609	-	0.7913	0.8611	0.8799	0.8589	0.8517	0.8568
	20	-	70	0.6774	-	0.8414	0.9231	0.9780	0.9726	0.9407	0.9642
CT	-	36	40	-	0.4919	0.7657	0.9053	0.9228	0.9178	0.9068	0.9129
	-	36	70	-	0.5895	0.8232	0.9175	0.9319	0.9290	0.9133	0.9181
	-	60	40	-	0.6726	0.8637	0.9390	0.9535	0.9520	0.9422	0.9467
	-	60	70	-	0.7846	0.9204	0.9548	0.9625	0.9626	0.9507	0.9576

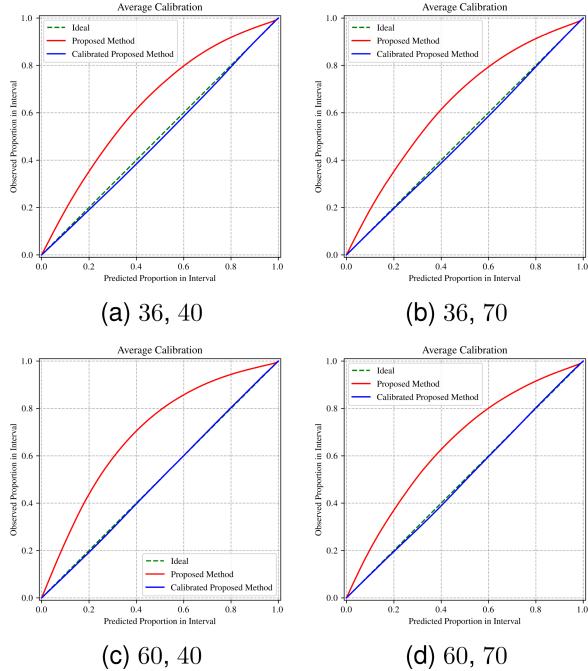


Fig. 9. Calibration plots of the proposed method for different configurations of the CT setup. Subcaptions specify the number of views and the SNR (dB).

and CT reconstruction problems, respectively. The reconstruction performance of POEM shows a decrease compared to PUMw/oBN due to using dropout after every convolutional layer, which is a strong form of regularization. Similarly, we observe that the reconstruction performance of POEM is slightly worse than that of POAM. The reconstruction performance of the proposed framework shows a decrease compared to POAM because of using dropout after every convolutional layer, which is a strong form of regularization. Comparing the proposed framework with POAM, the proposed framework shows an average SSIM decrease of 0.010 for the MRI reconstruction problem and 0.007 for the CT reconstruction problem. We observe a similar trend for the proposed framework and PUMw/oBN. On the other hand, the proposed framework achieves average SSIM gains of 0.010 and 0.006 for the MRI and CT reconstruction problems when compared to POEM, respectively. Similarly, the proposed framework surpasses the state-of-the-art deep unrolling method PUM. Due to space limitations, only representative visual results are presented in Fig. 10. Detailed visual results are provided in the Supplementary Material.

## V. DISCUSSION

Quantification of the epistemic uncertainty is crucial for learning-based image reconstruction methods, especially in safety-critical imaging applications, for quantifying the confidence on a reconstruction obtained using a model learned from available, potentially limited or unrepresentative training data. Our experimental results presented in Section IV showed that the epistemic uncertainty information provided by the proposed method exhibits the reducibility property of the epistemic uncertainty. Moreover, the epistemic uncertainty provided by the proposed framework can be used to assess how uncertain the learning-based image reconstruction method is and to detect cases where the input contains abnormal features not present in the training data.

For ill-posed inverse problems encountered in most imaging problems, inherent uncertainty on the target image for a given measurement vector is inevitable. Hence, it is essential to quantify the aleatoric uncertainty for imaging problems to capture the inherent randomness in the reconstruction task. Our experiments presented in Section IV demonstrated that the proposed framework is capable of capturing the aleatoric uncertainty in the sense that the aleatoric uncertainty provided by the proposed framework shows the regions where the reconstruction is expected to be challenging for the unrolled network. The aleatoric uncertainty provided by the proposed framework can be utilized to determine the possible errors in the reconstructed image and can be used as a mechanism to further assess the reliability of the reconstructed image. As a result, the aleatoric and epistemic uncertainties provided by the proposed framework would open the possibility of developing more accurate, robust, trustworthy, uncertainty-aware, learning-based image reconstruction and analysis methods. While the uncertainty estimates provided by our proposed methodology appear to reflect expected behavior of epistemic and aleatoric uncertainties, further analysis of the implications of the variational inference approximations used here would be beneficial.

The benefits of obtaining the epistemic and aleatoric uncertainty maps come with a price. Because the proposed framework requires feeding the measurement vector into the neural networks  $T$  times for inference, the inference time of the proposed framework increases by  $T$  times compared to the state-of-the-art deep unrolling method PUM. To shorten the inference time of the proposed framework, we can perform those  $T$  forward passes in parallel. Assuming that the GPU memory allows using a batch size of  $B$  in the inference stage, the proposed framework requires only  $\lfloor T/B \rfloor + 1$  forward passes for inference. If we

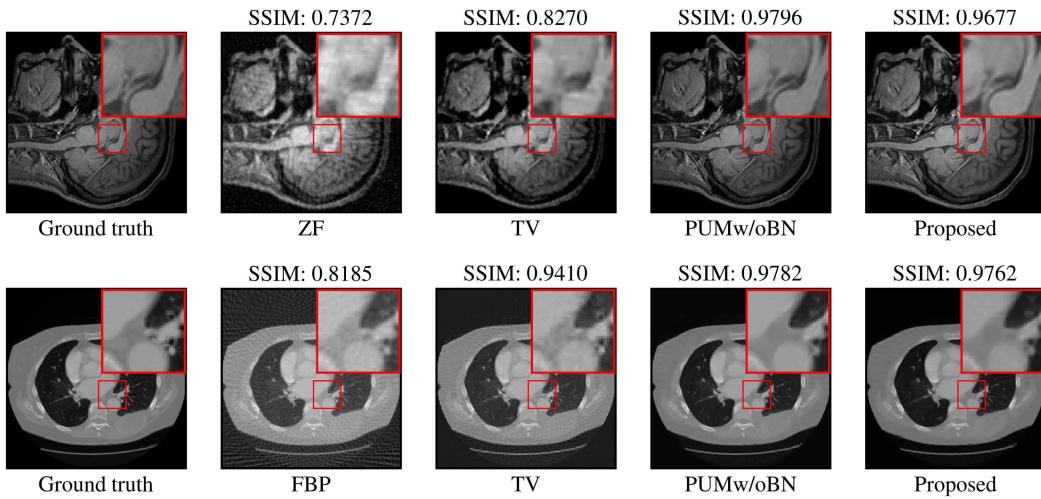


Fig. 10. Visual comparison of the image reconstruction performance of zero-filling (ZF) / filtered backprojection (FBP), total variation reconstruction (TV), state-of-the-art PGD-based deep unrolling method without batch normalization (PUMw/oBN), and the proposed method. Proposed method achieves comparable reconstruction performance to the state-of-the-art deep unrolling method PUMw/oBN, while providing uncertainty quantification.

have multiple GPUs, the inference time of the proposed framework can be further reduced. Hence, the proposed framework can achieve shorter inference times at the expense of using more computational power. Another way to shorten the inference time of the proposed framework is to decrease the number of parameters the proposed framework so that a larger batch size  $B$  can be used to parallelize the inference stage. To that end, we can design a variant of the proposed framework that uses a dual-head network. For the sake of brevity, we have not discussed this variant; however, a brief discussion on that variant is provided in the Supplementary Material.

## VI. CONCLUSION

In this article, we utilized the idea of deep unrolling and Bayesian neural networks to propose a learning-based image reconstruction framework that is capable of quantifying epistemic and aleatoric uncertainties while incorporating the imaging observation model into the reconstruction process. Our experimental results showed that the proposed framework provides epistemic and aleatoric uncertainty maps while providing a reconstruction performance comparable to the state-of-the-art deep unrolling methods. The proposed framework can be applied to a broad set of imaging problems and can be easily implemented in deep learning frameworks. We hope that the proposed framework and the provided discussion on epistemic and aleatoric uncertainties for imaging problems motivate further research on uncertainty characterization for imaging problems and on leveraging the uncertainty information for image reconstruction and analysis tasks.

## REFERENCES

- [1] I. Elbakri and J. Fessler, "Statistical image reconstruction for polyenergetic x-ray computed tomography," *IEEE Trans. Med. Imag.*, vol. 21, no. 2, pp. 89–99, Feb. 2002.
- [2] J. A. Fessler, "Model-based image reconstruction for MRI," *IEEE Signal Process. Mag.*, vol. 27, no. 4, pp. 81–89, Jul. 2010.
- [3] P. Sarder and A. Nehorai, "Deconvolution methods for 3-D fluorescence microscopy images," *IEEE Signal Process. Mag.*, vol. 23, no. 3, pp. 32–45, May 2006.
- [4] L. C. Potter, E. Ertin, J. T. Parker, and M. Cetin, "Sparsity and compressed sensing in radar imaging," *Proc. IEEE Proc.*, vol. 98, no. 6, pp. 1006–1020, Jun. 2010.
- [5] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, no. 1, pp. 89–97, 2004.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [7] N. Parikh and S. Boyd, "Proximal algorithms," *Foundations Trends Optim.*, vol. 1, no. 3, pp. 127–239, 2014.
- [8] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proc. Int. Conf. Int. Conf. Mach. Learn.*, 2010, pp. 399–406.
- [9] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-based deep learning architecture for inverse problems," *IEEE Trans. Med. Imag.*, vol. 38, no. 2, pp. 394–405, Feb. 2019.
- [10] J. Adler and O. Oktem, "Learned primal-dual reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1322–1332, Jun. 2018.
- [11] M. Borgerding, P. Schniter, and S. Rangan, "AMP-inspired deep networks for sparse linear inverse problems," *IEEE Trans. Signal Process.*, vol. 65, no. 16, pp. 4293–4308, Aug. 2017.
- [12] Y. Chun and J. A. Fessler, "Deep BCD-Net using identical encoding-decoding CNN structures for iterative image recovery," in *Proc. IEEE Image, Video, Multidimensional Signal Process. Workshop*, 2018, pp. 1–5.
- [13] S. Diamond, V. Sitzmann, F. Heide, and G. Wetzstein, "Unrolled optimization with deep priors," 2017, *arXiv:1705.08041*.
- [14] R. Liu, S. Cheng, L. Ma, X. Fan, and Z. Luo, "Deep proximal unrolling: Algorithmic framework, convergence analysis and applications," *IEEE Trans. Image Process.*, vol. 28, no. 10, pp. 5013–5026, Oct. 2019.
- [15] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-Net for compressive sensing MRI," in *Proc. Ad. Neural Inf. Process. Syst.*, 2016, vol. 29, pp. 10–18.
- [16] M. Mardani et al., "Neural proximal gradient descent for compressive imaging," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 9596–9606.
- [17] R. M. Neal, "Bayesian learning for neural networks," Ph.D. dissertation, Univ. Toronto, Toronto, ON, Canada, 1995. [Online]. Available: <https://www.cs.toronto.edu/radford/ftp/thesis.pdf>
- [18] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5580–5590.
- [19] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervention*, 2015, pp. 234–241.

[20] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.

[21] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 2796–2804.

[22] J. Adler and O. Oktem, "Deep posterior sampling: Uncertainty quantification for large scale inverse problems," in *Proc. Int. Conf. Med. Imag. with Deep Learn.*, 2019.

[23] J. M. Bardsley, "MCMC-based image reconstruction with uncertainty quantification," *SIAM J. Sci. Comput.*, vol. 34, no. 3, pp. A1316–A1332, 2012.

[24] V. Bohm, F. Lanusse, and U. Seljak, "Uncertainty quantification with generative models," in *Proc. NeurIPS Bayesian Deep Learn. Workshop*, 2019.

[25] X. Cai, M. Pereyra, and J. D. McEwen, "Uncertainty quantification for radio interferometric imaging—I. Proximal MCMC methods," *Monthly Notices Roy. Astronomical Soc.*, vol. 480, no. 3, pp. 4154–4169, 2018.

[26] X. Cai, M. Pereyra, and J. D. McEwen, "Uncertainty quantification for radio interferometric imaging: II. MAP estimation," *Monthly Notices Roy. Astronomical Soc.*, vol. 480, no. 3, pp. 4170–4182, 2018.

[27] J. M. Cochrane, M. Beveridge, and I. Drori, "Generalizing imaging through scattering media with uncertainty estimates," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. Workshops*, 2022, pp. 760–766.

[28] A. Dasgupta and Z. W. Di, "Uncertainty quantification for ptychography using normalizing flows," in *Proc. NeurIPS Mach. Learn. Phys. Sci. Workshop*, 2021.

[29] V. Edupuganti, M. Mardani, S. Vasanawala, and J. Pauly, "Uncertainty quantification in deep MRI reconstruction," *IEEE Trans. Med. Imag.*, vol. 40, no. 1, pp. 239–250, Jan. 2021.

[30] C. Ekmekci and M. Cetin, "What does your computational imaging algorithm not know?: A Plug-and-Play model quantifying model uncertainty," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops*, 2021, pp. 4018–4027.

[31] C. Ekmekci and M. Cetin, "Model-based Bayesian deep learning architecture for linear inverse problems in computational imaging," *Electron. Imag.*, vol. 2021, pp. 201-1–201-7, 2021.

[32] F. J. Herrmann, A. Siahkoohi, and G. Rizzuti, "Learned imaging with constraints and uncertainty quantification," in *Proc. NeurIPS Workshop Solving Inverse Problems Deep Netw.*, 2019.

[33] L. Hoffmann, I. Fortmeier, and C. Elster, "Uncertainty quantification by ensemble learning for computational optical form measurements," *Mach. Learn.: Sci. Technol.*, vol. 2, no. 3, 2021, Art. no. 035030.

[34] R. Kitichotkul, C. A. Metzler, F. Ong, and G. Wetzstein, "Suremap: Predicting uncertainty in CNN-based image reconstructions using Stein's unbiased risk estimate," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 1385–1389.

[35] R. Liu, S. Cheng, L. Tian, and J. Yi, "Deep spectral learning for label-free optical imaging oximetry with uncertainty quantification," *Light, Sci. Appl.*, vol. 8, pp. 1–13, 2019.

[36] A. Repetti, M. Pereyra, and Y. Wiaux, "Scalable Bayesian uncertainty quantification in imaging inverse problems via convex optimization," *SIAM J. Imag. Sci.*, vol. 12, no. 1, pp. 87–118, 2019.

[37] J. Schlemper et al., "Bayesian deep learning for accelerated MR image reconstruction," in *Proc. Int. Workshop Mach. Learn. Med. Image Reconstruction*, 2018, pp. 64–71.

[38] R. Shang, M. A. O'Brien, and G. P. Luke, "Deep-learning-driven reliable single-pixel imaging with uncertainty approximation," 2021, *arXiv:2107.11678*.

[39] A. Siahkoohi, G. Rizzuti, and F. Herrmann, "A deep-learning based Bayesian approach to seismic imaging and uncertainty quantification," in *Proc. EAGE 2020 Annu. Conf. Exhib. Online*, 2020, pp. 1–5.

[40] H. Sun and K. L. Bouman, "Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, pp. 2628–2637.

[41] R. Tanno et al., "Uncertainty quantification in deep learning for safer neuroimage enhancement," 2019, *arXiv:1907.13418*.

[42] F. Tonolini, J. Radford, A. Turpin, D. Faccio, and R. Murray-Smith, "Variational inference for computational imaging inverse problems," *J. Mach. Learn. Res.*, vol. 21, no. 179, pp. 1–46, 2020. [Online]. Available: <http://jmlr.org/papers/v21/20-151.html>

[43] Y. Xue, S. Cheng, Y. Li, and L. Tian, "Reliable deep-learning-based phase imaging with uncertainty quantification," *Optica*, vol. 6, no. 5, pp. 618–629, 2019.

[44] S. Bond-Taylor, A. Leach, Y. Long, and C. G. Willcocks, "Deep generative modelling: A comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 11, pp. 7327–7347, 2022.

[45] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*.

[46] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, vol. 70, pp. 214–223.

[47] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *Proc. Int. Conf. Learn. Representations*, 2014.

[48] I. Kobyzev, S. J. Prince, and M. A. Brubaker, "Normalizing flows: An introduction and review of current methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 11, pp. 3964–3979, Nov. 2021.

[49] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, "Normalizing flows for probabilistic modeling and inference," *J. Mach. Learn. Res.*, vol. 22, no. 57, pp. 1–64, 2021.

[50] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, pp. 491–503, 2018.

[51] W. Shi et al., "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, 2016, pp. 1874–1883.

[52] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 9446–9454.

[53] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. V. Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 6360–6376, Oct. 2022.

[54] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2013, pp. 945–948.

[55] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, "Hands-on Bayesian neural networks—A tutorial for deep learning users," *IEEE Comput. Intell. Mag.*, vol. 17, no. 2, pp. 29–48, May 2022.

[56] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.

[57] S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udluft, "Learning and policy search in stochastic dynamical systems with bayesian neural networks," in *Proc. Int. Conf. Learn. Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=H1f8S9ee>

[58] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.

[59] A. Kendall and R. Cipolla, "Modelling uncertainty in deep learning for camera relocation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 4762–4769.

[60] F. K. Gustafsson, M. Danelljan, and T. B. Schon, "Evaluating scalable Bayesian deep learning methods for robust computer vision," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 318–319.

[61] Y. Gal and Z. Ghahramani, "Bayesian convolutional neural networks with Bernoulli approximate variational inference," in *Proc. Int. Conf. Learn. Representations*, 2016.

[62] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 56, pp. 1929–1958, 2014.

[63] B. Lakshminarayanan et al., "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, vol. 30, pp. 6402–6413. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf>

[64] D. Hill, S. Williams, D. Hawkes, and S. Smith, "IXI dataset: IXI - information eXtraction from images project (EPSRC GR/S21533/02)," [Online]. Available: <https://brain-development.org/ixi-dataset/>

[65] A. A. A. Setio et al., "Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: The LUNA16 challenge," *Med. Image Anal.*, vol. 42, pp. 1–13, 2016.

[66] Y. Chung, I. Char, H. Guo, J. Schneider, and W. Neiswanger, "Uncertainty toolbox: An open-source library for assessing, visualizing, and improving uncertainty quantification," in *Proc. Uncertainty Robustness Deep Learn. Workshop Int. Conf. Mach. Learn.*, 2021.

[67] H. J. Yong, D. Huang, X. M. Hua, and L. Zhang, "Momentum batch normalization for deep learning with small batch size," in *Proc. Eur. Conf. Comput. Vis.*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds., 2020, pp. 224–240.