

Model-Based Bayesian Deep Learning Architecture for Linear Inverse Problems in Computational Imaging

Canberk Ekmekci¹, Mujdat Cetin^{1,2}

¹ Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY, USA.

² Goergen Institute for Data Science, University of Rochester, Rochester, NY, USA.

Abstract

We propose a neural network architecture combined with specific training and inference procedures for linear inverse problems arising in computational imaging to reconstruct the underlying image and to represent the uncertainty about the reconstruction. The proposed architecture is built from the model-based reconstruction perspective, which enforces data consistency and eliminates the artifacts in an alternating manner. The training and the inference procedures are based on performing approximate Bayesian analysis on the weights of the proposed network using a variational inference method. The proposed architecture with the associated inference procedure is capable of characterizing uncertainty while performing reconstruction with a model-based approach. We tested the proposed method on a simulated magnetic resonance imaging experiment. We showed that the proposed method achieved an adequate reconstruction capability and provided reliable uncertainty estimates in the sense that the regions having high uncertainty provided by the proposed method are likely to be the regions where reconstruction errors occur.

Introduction

The problem of reconstructing the underlying image can be cast as solving a linear inverse problem for several imaging modalities such as magnetic resonance imaging [1] and computed tomography [2]. Thus, linear inverse problems are at the foundation of the computational imaging.

Recently, neural network-based methods have become increasingly popular to solve the linear inverse problems arising in computational imaging (for a review see [3]). While one class of methods such as [4] tries to invert the forward model with a deep neural network to reconstruct the underlying image from the measured data, another class of methods such as [5] takes a slightly more conservative approach and aims to recover the latent image by formulating an optimization problem and replacing some part of the iterative reconstruction algorithm with a neural network.

Although these methods achieve state-of-the-art results in several computational imaging applications, the resulting image might experience unexpected instabilities [6]. The lack of uncertainty information about the reconstructed image severely limits the applicability of neural network-based methods in practice, especially in safety-critical applications. If a neural network-based reconstruction algorithm is able to provide uncertainty information about the reconstructed image, that can be leveraged to assess the quality of the reconstruction or to warn the practitioner,

who uses a fully automated pipeline performing reconstruction and analysis tasks simultaneously. Thus, being able to represent the uncertainty in neural network-based reconstruction methods is crucial for computational imaging problems.

To solve the problem of obtaining uncertainty information with a Bayesian perspective, we need to define probability distributions on the weights of a neural network, and to obtain the posterior distribution of the weights. These type of probabilistic models, called Bayesian neural networks [7], have attracted significant attention recently. Unfortunately, obtaining the posterior distribution of weights is not an easy task because of the large number of parameters and complex neural network models. Hence, different approximation methods (see [8] and the references therein) have been used in the literature to perform approximate Bayesian analysis for neural networks.

In this article, we introduce a model-based Bayesian deep learning architecture for computational imaging, which is built by using a model-based reconstruction approach. The training and inference procedures utilize Monte Carlo (MC) Dropout [9] to perform variational inference. Combined with specific training and inference procedures, the proposed architecture performs reconstruction and provides uncertainty information about the reconstructed image. The proposed architecture and the associated training and inference procedures are easy to implement in deep learning frameworks.

Recently, [10] has proposed a U-Net [11] based methodology combined with MC Dropout to obtain uncertainty information about the reconstruction for phase imaging. The differences between our approach and the one in [10] lie in that our architecture is built by taking a model-based reconstruction perspective, which can be perceived as a slightly more conservative approach compared to [10], and that we focus on a broad class of computational imaging problems, i.e., all computational imaging problems that can be written as a system of linear equations.

The rest of the paper is organized as follows: In the “Preliminaries” section, we review the basic material related to the formulation of the reconstruction problem, the generative model of the data, and the Bayesian approach to uncertainty estimation. In the “Proposed Method” section, we describe the proposed architecture with its training and inference procedures in detail. In the “Experiments” section, we present the experimental results we have obtained in a simulated magnetic resonance imaging scenario. Finally, the “Conclusion” section concludes the paper.

Preliminaries

In this section, we review some background on the classical statistical reconstruction approach, the generative model of

This work was partially supported by the National Science Foundation (NSF) under Grant CCF-1934962.

the data, and the Bayesian approach to uncertainty estimation for neural networks.

MAP Reconstruction

Consider the following setup

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

where $\mathbf{m} \in \mathbb{C}^S$ is the vector containing the measurements, $\mathbf{A} \in \mathbb{C}^{S \times D}$ is the operator that represents the transformation applied by the imaging system, $\mathbf{s} \in \mathbb{C}^D$ is the vectorized image, and $\mathbf{n} \in \mathbb{C}^S$ is the measurement noise in the system, which is assumed to be circularly-symmetric complex Gaussian noise with variance σ_{noise}^2 .

For an underdetermined system ($S < D$), recovering the image, \mathbf{s} , from the measurements, \mathbf{m} , becomes an ill-posed problem. One way to restrict the solution space and regularize the task is to use the prior knowledge about the image, \mathbf{s} . Then, the maximum *a posteriori* (MAP) estimate of the image, $\hat{\mathbf{s}}$, can be found by solving the following optimization problem

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s} \in \mathbb{C}^D} \left\{ \|\mathbf{A}\mathbf{s} - \mathbf{m}\|_2^2 + \beta \psi(\mathbf{s}) \right\}, \quad (2)$$

where the function $\psi: \mathbb{C}^D \rightarrow \mathbb{R}$ is the regularizer coming from the image prior, and the parameter β controls the balance between the data fidelity term and the regularizer. To find an equivalent problem involving only real vectors and matrices, we introduce two operators $\eta: \mathbb{C}^n \rightarrow \mathbb{R}^{2n}$ and $\kappa: \mathbb{C}^{n \times m} \rightarrow \mathbb{R}^{2n \times 2m}$ such that

$$\eta_n(\mathbf{x}) = [\Re(\mathbf{x}^\top) \quad \Im(\mathbf{x}^\top)]^\top, \quad \kappa_{n,m}(\mathbf{X}) = \begin{bmatrix} \Re(\mathbf{X}) & -\Im(\mathbf{X}) \\ \Im(\mathbf{X}) & \Re(\mathbf{X}) \end{bmatrix}, \quad (3)$$

where \Re and \Im compute the element-wise real and imaginary parts of a given vector or matrix, respectively. Then, using these two operators, the optimization problem in Eqn. (2) can be written as

$$\hat{\mathbf{s}} = \eta_D^{-1} \left(\arg \min_{\tilde{\mathbf{s}} \in \mathbb{R}^{2D}} \left\{ \|\tilde{\mathbf{A}}\tilde{\mathbf{s}} - \tilde{\mathbf{m}}\|_2^2 + \beta \tilde{\psi}(\tilde{\mathbf{s}}) \right\} \right), \quad (4)$$

where $\tilde{\mathbf{A}} := \kappa_{S,D}(\mathbf{A})$, $\tilde{\mathbf{s}} := \eta_D(\mathbf{s})$, $\tilde{\mathbf{m}} := \eta_S(\mathbf{m})$, and $\tilde{\psi}$ is the modified regularizer such that $\psi(\mathbf{s}) = \tilde{\psi}(\tilde{\mathbf{s}})$ for all $(\mathbf{s}, \tilde{\mathbf{s}})$ pairs. Therefore, finding the solution of the Eqn. (2) boils down to finding the solution of the optimization problem located inside the η_D^{-1} operator in Eqn. (4).

Assuming that the modified regularizer, $\tilde{\psi}$, is a convex, closed and proper function, we can use different splitting methods such as the proximal gradient method [12] or the alternating direction of method of multipliers (ADMM) [13] to solve the problem inside the η_D^{-1} operator in Eqn. (4) efficiently.

Generative Model of the Data

Suppose for a fixed operator \mathbf{A} and a noise variance σ_{noise}^2 , we have a training dataset, \mathcal{D}_{tr} , consisting of a collection of N measurement-target image pairs, i.e.,

$$\mathcal{D}_{tr} = \{(\tilde{\mathbf{m}}^{(i)}, \tilde{\mathbf{s}}^{(i)}) \mid \tilde{\mathbf{m}}^{(i)} = \eta_S(\mathbf{m}^{(i)}), \tilde{\mathbf{s}}^{(i)} = \eta_D(\mathbf{s}^{(i)}), i \in [N]\}. \quad (5)$$

We assume that the training dataset, \mathcal{D}_{tr} , consists of i.i.d. samples of measurements and target images drawn from the distribution

$$(\tilde{\mathbf{m}}^{(i)}, \tilde{\mathbf{s}}^{(i)}) \sim p(\tilde{\mathbf{m}})p(\tilde{\mathbf{s}}|\tilde{\mathbf{m}}), \quad \forall i \in \{1, \dots, N\},$$

where both $p(\tilde{\mathbf{m}})$ and $p(\tilde{\mathbf{s}}|\tilde{\mathbf{m}})$ are unknown distributions and could be only accessed empirically through \mathcal{D}_{tr} . This assumption is justifiable for computational imaging problems. Each measurement in the training dataset is obtained by using the operator \mathbf{A} and the noise variance σ_{noise}^2 , which can be thought as generating a sample from $p(\tilde{\mathbf{m}})$. Ideally, we want to use ground truth images as target images; however, we do not have an access to ground truth images in practice. Thus, the target image for the measurement is obtained by imaging the same object using a suppressed noise level and without data reduction so that the resulting image, which is often referred to as the reference image, is close to the ground truth image. This process can be thought as generating a sample from $p(\tilde{\mathbf{s}}|\tilde{\mathbf{m}})$.

We also assume that the test dataset,

$$\mathcal{D}_{te} = \{(\tilde{\mathbf{m}}_*^{(i)}, \tilde{\mathbf{s}}_*^{(i)}) \mid \tilde{\mathbf{m}}_*^{(i)} = \eta_S(\mathbf{m}_*^{(i)}), \tilde{\mathbf{s}}_*^{(i)} = \eta_D(\mathbf{s}_*^{(i)}), i \in [M]\}, \quad (6)$$

consists of i.i.d. samples of measurements and target images drawn from the same distribution that the training measurements and target images are drawn from.

Bayesian Approach to Uncertainty Estimation

Suppose we have a training dataset, \mathcal{D}_{tr} , and a test measurement, $\tilde{\mathbf{m}}_*$, from \mathcal{D}_{te} . For a regression model with a set of parameters γ , the predictive distribution is

$$p(\tilde{\mathbf{s}}_*|\tilde{\mathbf{m}}_*, \mathcal{D}_{tr}) = \int p(\tilde{\mathbf{s}}_*|\tilde{\mathbf{m}}_*, \gamma)p(\gamma|\mathcal{D}_{tr})d\gamma. \quad (7)$$

Throughout this paper, we refer to $p(\tilde{\mathbf{s}}_*|\tilde{\mathbf{m}}_*, \gamma)$ as the likelihood, and to $p(\gamma|\mathcal{D}_{tr})$ as the posterior distribution.

Unfortunately, obtaining the posterior distribution is not an easy task because of the large number of parameters and complex neural network architectures. Hence, approximation techniques such as variational inference methods [9] or Markov Chain Monte Carlo-based methods [7] are necessary.

After obtaining an approximation of the posterior distribution and defining the form of the likelihood, we can approximate the integral in Eqn. (7) using Monte Carlo integration to estimate the predictive distribution. Finally, we obtain the mean and the variance of the estimated predictive distribution with moment-matching. The resulting predictive variance can be used to characterize the uncertainty in the prediction.

Proposed Method

We define the following Gaussian likelihood for the reconstruction problem

$$p(\tilde{\mathbf{s}}|\tilde{\mathbf{m}}, \gamma) = N(\tilde{\mathbf{s}}|f_\omega(\tilde{\mathbf{m}}), \text{diag}(\sigma_\phi(\tilde{\mathbf{m}})^2)), \quad (8)$$

where $f_\omega: \mathbb{R}^{2S} \rightarrow \mathbb{R}^{2D}$ and $\sigma_\phi: \mathbb{R}^{2S} \rightarrow \mathbb{R}^{2D}$ are neural networks parametrized by sets of parameters ω and ϕ , respectively, and $\gamma = \omega \cup \phi$.

The architecture of the neural network f_ω is motivated by the model-based reconstruction approach. Assuming that the modified regularizer is a closed proper convex function, the update equation of the proximal gradient method to solve the optimization problem inside the η^{-1} operator in Eqn. (4) is

$$\tilde{\mathbf{s}}^{[k]} = \text{prox}_{\lambda\beta\tilde{\psi}} \left\{ \left(\mathbf{I} - 2\lambda\tilde{\mathbf{A}}^\top\tilde{\mathbf{A}} \right) \tilde{\mathbf{s}}^{[k-1]} + 2\lambda\tilde{\mathbf{A}}^\top\tilde{\mathbf{m}} \right\}, \quad (9)$$




Figure 1. Proposed model-based Bayesian deep learning architecture.




Figure 2. Neural network block R used in the proposed model-based Bayesian deep learning architecture.

where λ is the step size, and $\tilde{\mathbf{s}}^{[k]}$ is the reconstructed image at the k^{th} iteration. The benefit of using proximal gradient method over methods such as ADMM is that it does not require any matrix inversion in the update equation.

Similar to the idea in [5], we replace the proximal operator $\text{prox}_{\lambda\beta\tilde{\Psi}}$ with a neural network, and the update equation becomes

$$\begin{aligned}\tilde{\mathbf{z}}^{[k]} &= (\mathbf{I} - 2\lambda\tilde{\mathbf{A}}^T\tilde{\mathbf{A}})\tilde{\mathbf{s}}^{[k-1]} + 2\lambda\tilde{\mathbf{A}}^T\tilde{\mathbf{m}} \\ \tilde{\mathbf{s}}^{[k]} &= R\{\tilde{\mathbf{z}}^{[k]}, \omega\},\end{aligned}\quad (10)$$

where R is the neural network, and ω is the set of parameters of the neural network R . For a fixed number of iterations K , the series of updates correspond to a deep neural network, which is the desired function f_ω . Figure 1 illustrates the structure of f_ω in detail. Figure 2 depicts the components of the neural network R that contains P residual blocks consisting of a skip connection and C modules. Each module includes a convolutional layer followed by dropout, and each module except the last one includes a rectified linear unit (ReLU) activation function. Note that we use the same neural network R at every iteration, so the resulting neural network f_ω is consistent with the model-based reconstruction approach.

For computational imaging problems, the neural network σ_ϕ captures the inherent noise in the target images, which is sometimes referred to as the aleatoric uncertainty [14]. Thus, the choice of the σ_ϕ depends on the requirements of the application. If the true conditional distribution, $p(\tilde{\mathbf{s}}|\tilde{\mathbf{m}})$, is not invariant across the training and test datasets, we can simply use an arbitrary neural network that provides a non-negative output to explicitly model the data dependent uncertainty. On the other hand, if the true conditional distribution is invariant across the training and test datasets, we can focus on quantifying the uncertainty on the model parameters and set it to be a constant function $\sigma_\phi = \sigma_{\text{model}}$, where σ_{model} is a fixed model parameter and $\phi = \emptyset$. In the rest of this section, we assume that σ_ϕ is a neural network parametrized by a set of parameters ϕ to keep the generality.

Specifying the form of the neural networks f_ω and σ_ϕ completely specifies the form of the likelihood function. Next, we need to find an approximation of the posterior distribution, $p(\gamma|\mathcal{D}_{tr})$, which represents the uncertainty in the parameters of the regression model, γ . The uncertainty in the parameters of the

regression model, which is sometimes referred to as the epistemic uncertainty [14], reflects the uncertainty caused by the lack of training examples around the test measurement $\tilde{\mathbf{m}}_*$.

Monte Carlo Dropout [9] is a popular and scalable variational inference method because it has a simple implementation in deep learning frameworks, requires minimal changes on the classical neural network architecture, training and inference procedures, and provides reliable uncertainty estimates in several applications such as depth completion [8], and semantic segmentation [14]. In this work, we employ Monte Carlo Dropout to obtain an approximation of the true posterior distribution $p(\gamma|\mathcal{D}_{tr})$ by minimizing the Kullback-Leibler divergence between the true posterior distribution and the distribution $q_\alpha(\gamma)$ parametrized by the set of parameters α .

If the likelihood is defined as in Eqn. (8), and the distribution that we use to approximate the true posterior distribution, $q_\alpha(\gamma)$, is a Bernoulli variational distribution [15], i.e.,

$$q_\alpha(\gamma) = \prod_{i=1}^L p(z_i = 1)N(\gamma_i|\alpha_i, \sigma^2) + p(z_i = 0)N(\gamma_i|0, \sigma^2), \quad (11)$$

where σ is a sufficiently small constant, $\alpha = \{\alpha_i\}_{i=1}^L$, $\gamma = \{\gamma_i\}_{i=1}^L$, and $p(z_i = 1) = p$, then the optimal set of parameters, α^* , can be approximated by solving the following optimization problem [14]

$$\begin{aligned}\alpha^* &= \arg \min_{\alpha} \left\{ \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^{2D} [\log[\sigma_{\tilde{\phi}^{(i)}}(\tilde{\mathbf{m}}^{(i)})]_k \right. \\ &\quad \left. + \frac{([\tilde{\mathbf{s}}^{(i)}]_k - [f_{\tilde{\omega}^{(i)}}(\tilde{\mathbf{m}}^{(i)})]_k)^2}{2[\sigma_{\tilde{\phi}^{(i)}}(\tilde{\mathbf{m}}^{(i)})]_k^2}] + \frac{p}{2N} \sum_{i=1}^L \alpha_i^2 \right\},\end{aligned}\quad (12)$$

where $[\cdot]_k$ denotes the k^{th} element of a given vector, and $\tilde{\gamma}^{(i)} = \tilde{\omega}^{(i)} \cup \tilde{\phi}^{(i)}$ is a sample from the Bernoulli variational distribution, $q_\alpha(\gamma)$.

Interestingly, generating a sample from the Bernoulli variational distribution requires sampling a set of Bernoulli random variables $\{z_i\}_{i=1}^L$ and multiplying them with the parameters of the Bernoulli variational distribution, $\{\alpha_i\}_{i=1}^L$. This procedure resembles the dropout operation in the deep learning literature. Hence, solving the optimization problem in Eqn. (12) boils down to training two neural networks f_ω and σ_ϕ using the first term of the Eqn. (12) as a loss function with the weight decay parameter of $\frac{p}{2N}$

while the dropout is applied after every weight layer of the neural networks f_ω and σ_ϕ with the dropout rate of $1 - p$. The resulting weights of the neural networks after the training stage are the optimal parameters of the Bernoulli variational distribution, α^* .

After obtaining an approximation of the posterior distribution, $q_{\alpha^*}(\gamma) \approx p(\gamma|\mathcal{D}_{tr})$, we can find approximations of the mean and the variance of the predictive distribution by moment matching with T samples

$$\mathbb{E}[\tilde{\mathbf{s}}_*|\tilde{\mathbf{m}}_*, \mathcal{D}_{tr}] \approx \frac{1}{T} \sum_{t=1}^T f_{\tilde{\omega}^{(t)}}(\tilde{\mathbf{m}}_*), \quad (13)$$

$$\begin{aligned} \text{Var}[[\tilde{\mathbf{s}}_*]_k|\tilde{\mathbf{m}}_*, \mathcal{D}_{tr}] &\approx \frac{1}{T} \sum_{t=1}^T [\sigma_{\tilde{\phi}^{(t)}}(\tilde{\mathbf{m}}_*)]_k^2 + \frac{1}{T} \sum_{t=1}^T [f_{\tilde{\omega}^{(t)}}(\tilde{\mathbf{m}}_*)]_k^2 \\ &\quad - \left(\frac{1}{T} \sum_{t=1}^T f_{\tilde{\omega}^{(t)}}(\tilde{\mathbf{m}}_*) \right)_k^2. \end{aligned} \quad (14)$$

Because the original formulation of the problem involves complex target images, we can find the mean and the variance of the complex version of the predictive distribution as

$$\mathbb{E}[\mathbf{s}_*|\mathbf{m}_*, \mathcal{D}_{tr}] = \eta_D^{-1} (\mathbb{E}[\tilde{\mathbf{s}}_*|\tilde{\mathbf{m}}_*, \mathcal{D}_{tr}]) \quad (15)$$

$$\text{Var}[[\mathbf{s}_*]_k|\mathbf{m}_*, \mathcal{D}_{tr}] = \text{Var}[[\tilde{\mathbf{s}}_*]_k|\tilde{\mathbf{m}}_*, \mathcal{D}_{tr}] + \text{Var}[[\tilde{\mathbf{s}}_*]_{k+D}|\tilde{\mathbf{m}}_*, \mathcal{D}_{tr}] \quad (16)$$

for all $k \in \{1, 2, \dots, D\}$. Remarkably, the inference procedure also requires samples from the Bernoulli variational distribution. Thus, finding the predictive variance and the mean boils down to feeding the measurement, $\tilde{\mathbf{m}}_*$, into the neural networks f_ω and σ_ϕ T times while the dropout is enabled and performing the updates in Eqn. (15) and Eqn. (16).

Experiments and Results

The main advantage of the proposed method is that we can both perform reconstruction and obtain uncertainty information about the reconstruction. In this section, we evaluate the proposed method on a simulated magnetic resonance imaging experiment and focus on addressing the following questions:

- Can we use the uncertainty information to locate the potentially erroneous regions in the reconstructed image?
- Is the uncertainty information reliable in the sense that the region described by the predictive mean and the predictive variance contains the target pixel intensity values?
- How does the size of the training dataset N effect the uncertainty?

Setup

In magnetic resonance imaging, the measurement vector, \mathbf{m} , contains the Fourier coefficients of the ground truth image. In experiments, we used a single-coil model with a constant sensitivity map; therefore, the forward operator \mathbf{A} is simply a subsampled Fourier transform operator. The subsampling rate of the Fourier coefficients in the k-space determines the structure of the \mathbf{A} matrix and the acceleration rate of the data acquisition step. For example,

if 20% of the Fourier coefficients are collected, the acceleration rate of the acquisition becomes $5\times$.

For different acceleration rates ($3.33\times$, $5\times$, and $10\times$) and noise standard deviations (0.01, 0.05, 0.1, and 0.2), we constructed 12 different training dataset-test dataset pairs. To obtain the target images of the training and test datasets, we extracted $N = 7300$ 256×256 MRI images from the MRI data of 558 patients in the IXI Dataset¹ and $M = 268$ 256×256 MRI images from the MRI data of 20 patients in the IXI Dataset, respectively. The target images were normalized such that the pixel intensity values lie between 0 and 1. For each acceleration rate-noise standard deviation pair, i.e., for each setup configuration, the measurements of the training and test datasets were generated using the linear model given in Eqn. (1).

For the function that maps measurements to images f_ω , we use the neural network architecture introduced in the ‘‘Proposed Method’’ section with $K = 10$ iterations, $P = 1$ residual blocks and $C = 8$ modules. We used 32 filters in the first and the second modules, 64 filters in the third and the fourth modules, 128 filters in the fifth, sixth, and the seventh modules, and 2 filters in the final module. The kernel size of the filters was set to 3, and the stride and the padding size were set to 1. By the assumption that the true conditional distribution is invariant between the training and the test datasets, here we concentrate on quantifying the uncertainty on the parameters of the regression model. Thus, we set σ_ϕ to be a constant function with a fixed model parameter $\sigma_{model} = 0.001$, i.e. $\sigma_\phi = \sigma_{model}$, and $\phi = \emptyset$. In the case where the target images contain noise, or the true conditional distribution is not invariant across the training and test datasets, representing the aleatoric uncertainty becomes an important task for characterizing the overall uncertainty in the prediction. We leave the investigation of these cases for future work.

Uncertainty Map and Reconstruction Error

To investigate the relationship between the uncertainty map obtained by the proposed method and the reconstruction error, we tested the proposed model on test measurements and computed the absolute reconstruction errors and the uncertainty maps. Figure 3 illustrates the reconstruction error and the uncertainty map for a single test measurement.

As can be seen in Figure 3, the similarity between the uncertainty map and the absolute reconstruction error is remarkable. The uncertainty provided by the proposed method is high in the regions where the reconstruction error occurs such as the regions around the edges and small details. On the other hand, it is relatively low in the regions where the reconstruction error is negligibly small such as piecewise constant and smooth regions. We observed the same behavior for different setup configurations as well. Thus, we deduce that we can leverage the uncertainty information provided by the proposed method to locate the erroneous regions in a reconstructed image without the need of the reference image.

Reliability of the Uncertainty Information

To evaluate the reliability of the uncertainty information obtained by the proposed method, we trained the proposed model for each setup configuration using the corresponding training dataset.

¹<https://brain-development.org/ixi-dataset/>




Figure 3. The target image, the reconstructed image, the reconstruction error, and the uncertainty map for a single test measurement.

We performed $T = 100$ stochastic forward passes on each measurement in the test datasets to obtain the reconstructed complex image and the corresponding uncertainty map. Then, we counted the number of pixels whose target intensity values lie within three predictive standard deviations from the predictive mean (the reconstruction). Figure 4 shows the percentage of the number of pixels satisfying this condition for each reconstructed complex image in the test dataset for each setup configuration.

The results indicate that significant amount of target pixel intensity values lie within the region determined by the predictive mean and the three predictive standard deviations from the predictive mean. One can argue that these percentages can be made arbitrary high if the predictive variance is high across the image; however, this is not the case because in the previous subsection, we experimentally showed that the predictive variance is high in the erroneous regions and low in the regions where reconstruction error is negligible or not present. Thus, we can claim that the uncertainty information obtained by the proposed method is reliable in the sense that it correctly localizes the erroneous and non-erroneous regions, and the target pixel intensity values lie within the region determined by the uncertainty information.

Uncertainty Map and Training Dataset Size

To assess the effect of the number of training examples in the dataset on the uncertainty, we trained four different versions of the proposed model with $N = 10$, $N = 50$, $N = 100$, and $N = 7000$ training examples. We performed $T = 100$ stochastic forward passes on each test example to obtain the reconstructed images and the corresponding uncertainty maps. Figure 5 shows the reconstructed images and the corresponding uncertainty maps obtained by the aforementioned models.

For the model that was trained on $N = 10$ examples, the model uncertainty obtained by the proposed method is severely high, and it suggests that the number of training examples must be increased to obtain a more confident reconstruction. For the second and the third models trained with $N = 50$ and $N = 100$




Figure 4. The percentage of the number of pixels whose target intensity values lie within three predictive standard deviations from the predictive mean for each reconstructed complex image in the test dataset for each setup configuration.

examples, the overall uncertainty in the prediction continued to decrease. For the models trained with $N = 100$ and $N = 7000$ examples, the overall decrease in the uncertainty levels was not as drastic as the first two cases, but the uncertainty levels kept decreasing in some of the smooth regions. Hence, the model uncertainty maps in Figure 5 suggest that the model would get more confident as we increase the size of the training set.

One can confirm this by investigating the differences between the reconstructions obtained by the four models in Figure 5. From left to right, the values of SSIM index are 0.899, 0.922, 0.929, and 0.932. Considering the contrast, the artifacts and the details, the best and the most confident reconstruction was obtained with the model trained with $N = 7000$ examples. Therefore, our results show that the proposed model is successful at capturing the uncertainty information depending on the size of the dataset.

Conclusion

We proposed a model-based Bayesian deep learning architecture for computational imaging and its specific training and inference procedures. The proposed architecture complies with the model-based reconstruction approach, and approximate Bayesian inference is performed using the MC Dropout technique. The proposed method is simple to implement in different deep learning frameworks and has a fast inference procedure. We tested the proposed method in magnetic resonance imaging simulation under a variety of setup configurations. We experimentally showed that the uncertainty information provided by the proposed method is reliable, and we can leverage it to detect the erroneous regions in the reconstructed image without the need of the reference image.

References

- [1] Jeffrey A. Fessler, Optimization Methods for Magnetic Resonance Image Reconstruction: Key Models and Optimization Algorithms, IEEE Signal Processing Magazine, vol. 37, no. 1, pp. 33–40, 2020.
- [2] Alireza Entezari, Masih Nilchian, and Michael Unser, A Box Spline Calculus for the Discretization of Computed Tomography Reconstruction Problems, IEEE Transactions on Medical Imaging, vol. 31, no. 8, pp. 1532–1541, 2012.
- [3] George Barbastathis, Aydogan Ozcan, and Guohai Situ, On the use

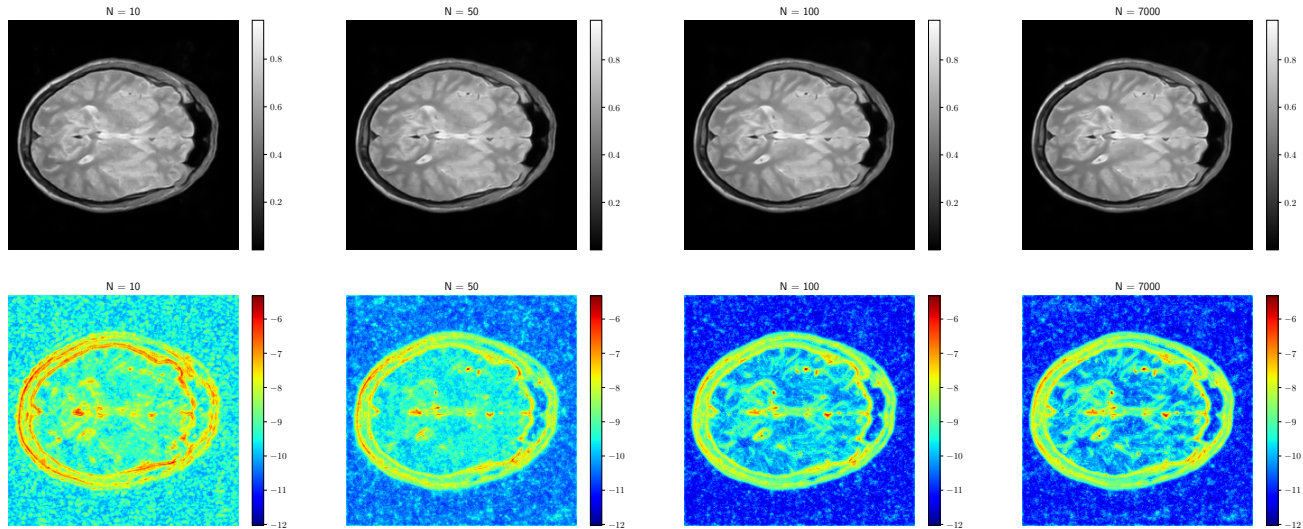


Figure 5. Effect of the number of training examples N in the training dataset on the uncertainty information. The first row shows the reconstructed images. The second row shows the corresponding predictive variance obtained by the proposed method for each reconstruction in log scale. The results in the first, second, third, and the fourth columns were obtained by using $N = 10$, $N = 50$, $N = 100$, and $N = 7000$ training examples, respectively.

of deep learning for computational imaging, *Optica*, vol. 6, no. 8, pp. 921–943, 2019.

- [4] Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser, Deep Convolutional Neural Network for Inverse Problems in Imaging, *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [5] Morteza Mardani, Sun Qingyun, Shreyas Vasawanala, Vardan Papayan, Hatf Monajemi, John Pauly, and David Donoho, Neural Proximal Gradient Descent for Compressive Imaging, *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 9596–9606, 2018.
- [6] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C. Hansen, On instabilities of deep learning in image reconstruction and the potential costs of AI, *Proceedings of the National Academy of Sciences*, 2020.
- [7] Radford M. Neal, Bayesian learning for neural networks, Ph.D. thesis, University of Toronto, 1995.
- [8] Fredrik K. Gustafsson, Martin Danelljan, and Thomas B. Schon, Evaluating scalable Bayesian deep learning methods for robust computer vision, *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 318–319, 2020.
- [9] Yarin Gal, and Zoubin Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, *International Conference on Machine Learning*, pp. 1050–1059, 2016.
- [10] Yujia Xue, Shiyi Cheng, Yunzhe Li, and Lei Tian, Reliable deep-learning-based phase imaging with uncertainty quantification, *Optica*, vol. 6, pp. 618–629, 2019.
- [11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, U-Net: Convolutional networks for biomedical image segmentation, *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, 2015.
- [12] Neal Parikh, and Stephen Boyd, Proximal algorithms, *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.
- [13] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein, Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers, *Foundations and Trends*

in Machine Learning, vol. 3, no. 1, pp. 1–122, 2010.

- [14] Alex Kendall, and Yarin Gal, What uncertainties do we need in Bayesian deep learning for computer vision?, *Advances in Neural Information Processing Systems*, pp. 5574–5584, 2017.
- [15] Yarin Gal, and Zoubin Ghahramani, Bayesian Convolutional Neural Networks with Bernoulli Approximate Variational Inference, *4th International Conference on Learning Representations (ICLR) workshop track*, 2016.

Author Biography

Canberk Ekmekci is currently pursuing the Ph.D. degree at the University of Rochester, Rochester, NY, under the supervision of Prof. Mujdat Cetin. His research interests include machine learning, convex optimization, and linear inverse problems arising in computational imaging.

Mujdat Cetin is a Professor of Electrical and Computer Engineering and the Director of the Goergen Institute for Data Science at the University of Rochester. His research interests include computational imaging, bioimage analysis, and brain-computer/machine interfaces. Prof. Cetin is a Fellow of the IEEE and has served as a member of the IEEE Signal Processing Society Technical Directions Board and as the Chair of the IEEE Computational Imaging Technical Committee. He is currently a Senior Area Editor for the *IEEE Transactions on Computational Imaging* and the *IEEE Transactions on Image Processing*. Prof. Cetin has received several awards including the IEEE Signal Processing Society Best Paper Award; the EURASIP/Elsevier Signal Processing Best Paper Award; the IET Radar, Sonar and Navigation Premium Award; and the Turkish Academy of Sciences Distinguished Young Scientist Award. He was the Technical Program Co-chair for the IEEE Image, Video, and Multidimensional Signal Processing (IVMSP) Workshop in 2016; for the International Conference on Information Fusion in 2016 and 2013; for the International Conference on Pattern Recognition (ICPR) in 2010; and for the IEEE Conference on Signal Processing, Communications, and their Applications in 2006.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

