

# Uncertainty Estimation for Computed Tomography with a Linearised Deep Image Prior

Anonymous authors

Paper under double-blind review

## Abstract

Existing deep-learning based tomographic image reconstruction methods do not provide accurate estimates of reconstruction uncertainty, hindering their real-world deployment. This paper develops a method, termed as the linearised deep image prior (DIP), to estimate the uncertainty associated with reconstructions produced by the DIP with total variation regularisation (TV). Specifically, we endow the DIP with conjugate Gaussian-linear model type error-bars computed from a local linearisation of the neural network around its optimised parameters. To preserve conjugacy, we approximate the TV regulariser with a Gaussian surrogate. This approach provides pixel-wise uncertainty estimates and a marginal likelihood objective for hyperparameter optimisation. We demonstrate the method on synthetic data and real-measured high-resolution 2D  $\mu$ CT data, and show that it provides superior calibration of uncertainty estimates relative to previous probabilistic formulations of the DIP. Our code is available at [https://github.com/anonymooseBayesDIP/bayes\\_dip](https://github.com/anonymooseBayesDIP/bayes_dip).

## 1 Introduction

Inverse problems in imaging aim to recover an unknown image  $x \in \mathbb{R}^{d_x}$  from the noisy measurement  $y \in \mathbb{R}^{d_y}$

$$y = Ax + \eta, \quad (1)$$

where  $A \in \mathbb{R}^{d_y \times d_x}$  is a linear forward operator, and  $\eta$  i.i.d. noise. We assume Gaussian noise  $\eta \sim \mathcal{N}(0, \sigma_y^2 I)$ . Many tomographic reconstruction problems take this form, e.g. computed tomography (CT). Due to the inherent ill-posedness of the reconstruction task, e.g.  $d_y \ll d_x$ , suitable regularisation, or prior specification, is crucial for the successful recovery of  $x$  (Tikhonov & Arsenin, 1977; Engl et al., 1996; Ito & Jin, 2014).

In recent years, deep-learning based approaches have achieved outstanding performance on a wide variety of tomographic problems (Arridge et al., 2019; Ongie et al., 2020; Wang et al., 2020). Most deep learning methods are supervised; they rely on large volumes of paired training data. Alas, these often fail to generalise out-of-distribution (Antun et al., 2020); small deviations from the distribution of the training data can lead to severe reconstruction artefacts. Pathologies of this sort motivate the need for both unsupervised deep learning methods—not dependent on training data and thus mitigating hallucinatory artefacts (Bora et al., 2017; Heckel & Hand, 2019; Tölle et al., 2021)—and uncertainty quantification (Kompa et al., 2021; Vasconcelos et al., 2022)—which can inform us about unreliability in reconstructions.

We focus on the deep image prior (DIP), perhaps the most widely adopted unsupervised deep learning approach (Ulyanov et al., 2018). DIP regularises the reconstructed image  $\hat{x}$  by reparametrising it as the output of a deep convolutional neural network (CNN). It does not require paired training data, relying solely




Figure 1: X-ray reconstruction ( $501 \times 501$  px<sup>2</sup>) of a walnut (left), the absolute error of its CT reconstruction (top) and pixel-wise uncertainty (bottom).

on the structural biases induced by the CNN architecture. The DIP has proven effective on tasks ranging from denoising and deblurring to more challenging tomographic reconstructions (Liu et al., 2019; Baguer et al., 2020; Knopp & Grosser, 2022; Darestani & Heckel, 2021; Gong et al., 2019; Cui et al., 2021; Barutcu et al., 2022). Nonetheless, the DIP only provides point reconstructions without uncertainty estimates.

In this work, we equip DIP reconstructions with reliable uncertainty estimates. Distinctly from previous probabilistic formulations of the DIP (Cheng et al., 2019; Tölle et al., 2021), we only estimate the uncertainty associated with a specific reconstruction, instead of trying to characterise a full posterior distribution over all candidate images. We achieve this by computing Gaussian-linear model type error-bars for a local linearisation of the DIP around its optimised reconstruction (Mackay, 1992; Khan et al., 2019; Immer et al., 2021b). Henceforth, we will refer to our method as *linearised DIP*. Linearised approaches have recently been shown to provide state-of-the-art uncertainty estimates for supervised deep learning models (Daxberger et al., 2021b). Unfortunately, the total variation (TV) regulariser, ubiquitous in CT reconstruction, makes inference in the linearised DIP intractable and it does not lend itself to standard Laplace (i.e. local Gaussian) approximations (Helin et al., 2022). We tackle this issue by using the predictive complexity prior (PredCP) framework of Nalisnick et al. (2021) to construct covariance kernels that induce properties similar to those of the TV prior while preserving Gaussian-linear conjugacy. Finally, we discuss a number of computational techniques that allow us to scale the proposed method to large DIP networks and high-resolution 2D reconstructions.

We demonstrate our approach on high-resolution CT reconstructions of real-measured 2D  $\mu$ CT projection data, cf. fig. 1. Empirically, our method’s pixel-wise uncertainty estimates predict reconstruction errors more accurately than existing approaches to uncertainty estimation with the DIP. This is not at the expense of accuracy in reconstruction: the reconstruction obtained using the standard regularised DIP method (Baguer et al., 2020) is preserved as the predictive mean, ensuring compatibility with advancements in DIP research.

The contributions of this work can be summarised as follows.

- We propose a novel approach to bestow reconstructions from the TV-regularised-DIP with uncertainty estimates. Specifically, we construct a local linear model by linearising the DIP around its optimised reconstruction and provide this model’s error-bars as a surrogate for those of the DIP.
- We detail an efficient implementation of our method, scaling up to real-measured high-resolution  $\mu$ CT data. In this setting, our method provides by far more accurate uncertainty estimation than previous probabilistic formulations of the DIP.

The rest of this paper is organised as follows. Section 2 introduces preliminaries necessary for our subsequent discussion of the linearised DIP. Section 3 discusses how to design a tractable Gaussian prior based on the TV semi-norm for Bayesian inversion of ill-posed imaging problems. Section 4 and section 5 present the linearised DIP along with its efficient implementation. Section 6 reports our experimental investigations carried out on synthetic and real-measured high-resolution  $\mu$ CT data. Section 7 concludes the article. The detailed derivations and additional experimental results are given in the supplementary material (SM).

Since this paper’s first appearance, the proposed method was used by Barbano et al. (2022) to actively select X-ray scanning angles, resulting in a 30% reduction in angles needed to obtain a given reconstruction PSNR, and extended by Antoran et al. (2023), scaling it to larger problems by drawing samples with SGD.

## 2 Preliminaries

### 2.1 Total variation regularisation

The imaging problem given in eq. (1) admits multiple solutions consistent with the observation  $y$ . Thus, regularisation is needed for stable reconstruction. Total variation (TV) is perhaps the most well established regulariser (Rudin et al., 1992; Chambolle et al., 2010). The anisotropic TV semi-norm of an image vector  $x \in \mathbb{R}^{d_x}$  imposes an  $L^1$  constraint on image gradients,

$$\text{TV}(x) = \sum_{i,j} |X_{i,j} - X_{i+1,j}| + \sum_{i,j} |X_{i,j} - X_{i,j+1}|, \quad (2)$$

where  $X \in \mathbb{R}^{h \times w}$  denotes the vector  $x$  reshaped into an image of height  $h$  by width  $w$ , and  $d_x = h \cdot w$ . This leads to the regularised reconstruction formulation

$$\hat{x} \in \operatorname{argmin}_{x \in \mathbb{R}^{d_x}} \mathcal{L}(x) \quad \text{with} \quad \mathcal{L}(x) := \|Ax - y\|_2^2 + \lambda \text{TV}(x), \quad (3)$$

where the hyperparameter  $\lambda > 0$  determines the strength of the regularisation relative to the fit term.

## 2.2 Bayesian inference for inverse problems

The probabilistic framework provides a consistent approach to uncertainty estimation in ill-posed imaging problems (Kaipio & Somersalo, 2005; Stuart, 2010; Seeger & Nickisch, 2011). In this framework, the image to be reconstructed is treated as a random variable. Instead of finding a single best reconstruction  $\hat{x}$ , we aim to find a posterior distribution  $p(x|y)$  that scores every candidate image  $x \in \mathbb{R}^{d_x}$  according to its agreement with our observations  $y$  and prior beliefs  $p(x)$ . Under this view, the regularised objective in eq. (3) can be understood as the negative log of an unnormalised posterior, i.e.  $p(x|y) \propto \exp(-\mathcal{L}(x))$ , and  $\hat{x}$  as its mode, i.e. the *maximum a posteriori* estimate. Specifically, the least squares reconstruction loss corresponds to the negative log of a Gaussian likelihood  $p(y|x) = \mathcal{N}(y; Ax, I)$  and the TV regulariser to the negative log of a prior density over reconstructions  $p(x) \propto \exp(-\lambda \text{TV}(x))$ .

The aforementioned posterior is obtained by updating the prior over images with the likelihood as

$$p(x|y) = p(y)^{-1} p(y|x) p(x), \quad (4)$$

for  $p(y) = \int p(y|x) p(x) dx$  the normalising constant, also known as the marginal likelihood (MLL). This latter quantity provides an objective for optimising hyperparameters, such as the regularisation strength  $\lambda$ . The discrepancy among plausible reconstructions in the posterior indicates uncertainty.

Our work partially departs from this framework in that it *solely concerns itself with characterising plausible reconstructions around the mode  $\hat{x}$*  (Mackay, 1992). This has two key advantages, i) *tractability*: the likelihood induced by NN reconstructions is strongly multi-modal, and both analytically and computationally intractable. On the other hand, the posterior for our local model is quadratic; ii) *interpretability*: even if we could obtain the full posterior, downstream stakeholders without expertise in probability are likely to have little use for it. A single reconstruction together with its pixel-wise uncertainty may be more interpretable to end-users.

## 2.3 The Deep Image Prior (DIP)

The DIP (Ulyanov et al., 2018; 2020) reparametrises the reconstructed image as the output of a CNN  $x(\theta)$  with learnable parameters  $\theta \in \mathbb{R}^{d_\theta}$  and a fixed input, which we have omitted from our notation for clarity. We overload  $x$  to refer to a point in  $\mathbb{R}^{d_x}$  and the injection  $x : \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}^{d_x}$  given by the DIP neural network. The DIP can be seen as a reparametrisation that provides a favourable loss landscape for optimisation (Shi et al., 2022). Penalising the TV of the DIP’s output avoids the need for early stopping and improves reconstruction fidelity (Liu et al., 2019; Baguer et al., 2020). The resulting optimisation problem is

$$\hat{\theta} \in \operatorname{argmin}_{\theta \in \mathbb{R}^{d_\theta}} \|Ax(\theta) - y\|_2^2 + \lambda \text{TV}(x(\theta)), \quad (5)$$

and the recovered image is given by  $\hat{x} = x(\hat{\theta})$ . U-Net is the standard choice of CNN architecture (Ronneberger et al., 2015). Although the parameters  $\theta$  must be optimised separately for each new measurement  $y$ , we follow (Barbano et al., 2021b; Knopp & Grosser, 2022) in reducing the cost of this with task-agnostic pretraining.

## 2.4 Bayesian inference with linearised neural networks

Adopting the DIP parametrisation of the reconstructed image, as in section 2.3, makes the Bayesian posterior in eq. (4) intractable. Instead, the present work only characterises the uncertainty associated with a specific regularised reconstruction  $\hat{x}$ , obtained per eq. (5). To this end, we take a first-order Taylor expansion of the CNN  $x(\theta)$  around its optimised parameters  $\hat{\theta}$  (Mackay, 1992; Khan et al., 2019; Immer et al., 2021b),

$$h(\theta) := x(\hat{\theta}) + J(\theta - \hat{\theta}), \quad (6)$$




Figure 2: Samples from priors. From left to right. Plot 1 shows a histogram of the average sample TV reporting an overlap between the TV and TV-PredCP priors. The factorised Gaussian prior results in larger TV values. Plot 2 shows an analogous histogram using samples from the linearised DIP (lin.-DIP) fitted to a KMNIST image, where the hyperparameters ( $\ell$ ,  $\sigma^2$ ) have been optimised both with and without the TV-PredCP term. Plots 3-5 show samples from the TV, TV-PredCP, and factorised Gaussian priors proposed in section 3, drawn using Hamiltonian Monte Carlo (HMC). Qualitatively, the TV prior produces samples with more correlated nearby pixel values than the factorised prior. The TV-PredCP prior captures this effect and produces even smoother samples, likely due to the presence of longer range correlation in the Matern-1/2 covariance. Plots 6-7 show prior samples from the linearised-DIP, which produces samples containing the structure of the KMNIST image used to train the network. The TV-PredCP term in the DIP hyperparameter optimisation to lead to smoother samples with less artefacts.

where  $J := \frac{\partial x(\theta)}{\partial \theta}|_{\theta=\hat{\theta}} \in \mathbb{R}^{d_x \times d_\theta}$  is the Jacobian of the CNN function with respect to its parameters evaluated at  $\hat{\theta}$ . We use this tangent linear model to obtain error-bars for the DIP reconstruction  $x(\hat{\theta})$ . When the observation noise is Gaussian, as is the case for the inverse problem formulation in eq. (1), and a Gaussian prior is placed on  $\theta$ , we recover the conjugate setting; the posterior distribution over the linearised model’s reconstructions is the Gaussian  $\mathcal{N}(x; \hat{x}, \Sigma_{x|y})$ , where we have fixed the mean to be the DIP reconstruction. In this conjugate setting, the marginal likelihood of the linearised model can be used to tune hyperparameters (Mackay, 1992; Immer et al., 2021a; Antorán et al., 2022; Antorán et al., 2022).

Computing both the posterior covariance  $\Sigma_{x|y}$  naively and the marginal likelihood has cost  $\mathcal{O}(d_\theta^3)$ . For large U-Nets, this is intractable (Daxberger et al., 2021b). In section 4 and section 5, we derive a dual approach with a cost  $\mathcal{O}(d_y^3)$  and detail an efficient implementation. Furthermore, when using the (non-quadratic) TV regulariser, conjugacy is lost. As we will see in section 3, the TV regulariser does not admit a Laplace (quadratic) approximation. Section 3 and section 4 are dedicated to addressing this issue.

### 3 The total variation as a conditionally Gaussian prior

The regularised reconstruction objective presented in eq. (3) can be interpreted as the negative log of an unnormalised posterior over reconstructions. In this context, the TV regulariser corresponds to the prior

$$p(x) = Z_\lambda^{-1} \exp(-\lambda \text{TV}(x)), \quad (7)$$

where  $Z_\lambda = \int \exp(-\lambda \text{TV}(x)) dx$  is its normalisation constant (the prior is improper, since constant vectors are in the null space of the derivative operator). Working with this prior is computationally intractable because  $Z_\lambda$  does not admit a closed form. The Laplace method, which consists of a locally quadratic approximation, does not solve the issue because the second derivative of the TV regulariser is zero everywhere it is defined.

As an alternative to enforce local smoothness in the reconstruction, we construct a Gaussian prior  $\mathcal{N}(x; 0, \Sigma_{xx})$  with covariance  $\Sigma_{xx} \in \mathbb{R}^{d_x \times d_x}$  given by the Matern-1/2 kernel

$$[\Sigma_{xx}]_{ij, i' j'} = \sigma^2 \exp\left(\frac{-d(i - i', j - j')}{\ell}\right), \quad (8)$$

where  $i, j$  index the spatial locations of pixels of  $x$ , as in eq. (2), and  $d(a, b) = \sqrt{a^2 + b^2}$ . The hyperparameter  $\sigma^2 \in \mathbb{R}^+$  informs the pixel amplitude while the lengthscale parameter  $\ell \in \mathbb{R}^+$  determines the correlation strength between nearby pixels. The expected TV associated with our Gaussian prior is

$$\kappa := \mathbb{E}_{x \sim \mathcal{N}(0, \Sigma_{xx})}[\text{TV}(x)] = c\sigma\sqrt{1 - \exp(-\ell^{-1})}, \quad (9)$$

with  $c$  a constant. See appendix B.3 for a derivation. Here and below, we may omit the dependence of  $\kappa$  on  $(\ell, \sigma^2)$  from our notation for clarity. For fixed pixel amplitude  $\sigma^2$ , the expected reconstruction TV  $\kappa$  is a bijection of the lengthscale  $\ell$ . We leverage this fact within the PredCP framework of Nalisnick et al. (2021) to construct a prior over  $\ell$  that will favour reconstructions with low expected TV

$$p(\ell) = \text{Exp}(\kappa) |\partial\kappa/\partial\ell|. \quad (10)$$

The resulting hierarchical prior over images, given by

$$x|\ell \sim \mathcal{N}(0, \Sigma_{xx}), \quad \ell \sim \text{Exp}(\kappa) |\partial\kappa/\partial\ell|, \quad (11)$$

is Gaussian for fixed  $\ell$ , and thus the prior is conditionally conjugate to Gaussian-linear likelihoods. Figure 2 shows agreement between samples, drawn with Hamiltonian Monte Carlo, from the described TV-PredCP prior and the intractable TV prior, both qualitatively and in terms of distribution over image TV.

## 4 The linearised DIP

In this section, we build a probabilistic model that aims to characterise posterior reconstructions around  $\hat{\theta}$ , a mode of the regularised DIP objective, which we assume to have obtained using eq. (5). Section 4.1 describes the construction of a linearised surrogate for the DIP reconstruction. Section 4.2 describes how to compute the surrogate model’s error-bars and use them to augment the DIP reconstruction. Section 4.3 discusses how we include the effects of TV regularisation into the surrogate model. Finally, in section 4.4, we describe a strategy to choose the surrogate model’s prior hyperparameters using a marginal likelihood objective.

### 4.1 From a prior over parameters to a prior over images

After training the DIP to an optimal TV-regularised setting  $\hat{x} = x(\hat{\theta})$  using eq. (5), we linearise the network around  $\hat{\theta}$  by applying eq. (6), and obtain the affine-in- $\theta$  function  $h(\theta)$ . The error-bars obtained from Bayesian inference with  $h(\theta)$  will tell us about the uncertainty in  $\hat{x}$ . To this end, consider the hierarchical model,

$$y|\theta \sim \mathcal{N}(A h(\theta), \sigma_y^2 I), \quad \theta|\ell \sim \mathcal{N}(0, \Sigma_{\theta\theta}(\ell)), \quad \ell \sim p(\ell) \quad \text{with} \quad h(\theta) := x(\hat{\theta}) + J(\theta - \hat{\theta}), \quad (12)$$

where we have placed a Gaussian prior over the parameters  $\theta$  that, in turn, depends on the lengthscale  $\ell$ . Conditional on a value of  $\ell$ , this is a conjugate Gaussian-linear model and thus the posterior distribution over  $\theta$  has a closed Gaussian form. The choice of distribution over the lengthscale will allow us to incorporate TV constraints into the computed error-bars, cf. section 4.3. We have introduced the noise variance  $\sigma_y^2$  as an additional hyperparameter which we will learn using the marginal likelihood (details in section 4.4).

To provide intuition about the linearised model, we push samples from  $\theta \sim \mathcal{N}(\theta; 0, \Sigma_{\theta\theta})$ , through  $h$ . The resulting reconstruction samples are drawn from a Gaussian distribution with covariance  $\Sigma_{xx} \in \mathbb{R}^{d_x \times d_x}$  given by  $J \Sigma_{\theta\theta} J^\top$  and are shown in fig. 2. Here, the Jacobian  $J$  introduces structure from the NN function around the linearisation point  $\hat{\theta}$ . This introduces features from the KMNIST character that the DIP was trained on.

### 4.2 Efficient posterior predictive computation

We augment the DIP reconstruction  $\hat{x}$  with Gaussian predictive error-bars computed with the linearised model  $h$  described in eq. (12), yielding  $\mathcal{N}(x; \hat{x}, \Sigma_{x|y})$ . The posterior covariance  $\Sigma_{x|y}$  is given by

$$\Sigma_{x|y} = J(\sigma_y^{-2} J^\top A^\top A J + \Sigma_{\theta\theta}^{-1})^{-1} J^\top = \Sigma_{xx} - \Sigma_{xy} \Sigma_{yy}^{-1} \Sigma_{xy}^\top, \quad (13)$$

which is derived in appendix C. Here,  $\Sigma_{xx} = J \Sigma_{\theta\theta} J^\top$ ,  $\Sigma_{xy} = \Sigma_{xx} A^\top$  and  $\Sigma_{yy} = A \Sigma_{xx} A^\top + \sigma_y^2 I$ . The constant-in- $\theta$  terms in  $h$  do not affect the uncertainty estimates, and thus the error-bars match those of the simple linear model  $J\theta$ . Importantly, eq. (13) depends on the inverse of the observation space covariance  $\Sigma_{yy}^{-1}$ , as opposed to the covariance over reconstructions, or parameters. Equation (13) scales as  $\mathcal{O}(d_x d_y^2)$  as opposed to  $\mathcal{O}(d_x^3)$  or  $\mathcal{O}(d_\theta^3)$  for the more-standard-in-the-literature output (reconstruction) space or parameter space approaches, respectively (Immer et al., 2021b; Daxberger et al., 2021a).

### 4.3 Incorporating TV-smoothness into our model as a prior

We aim to impose constraints on  $h$ 's error-bars, such that our model only considers low TV reconstructions as plausible. For this, we place a block-diagonal Matern-1/2 covariance Gaussian prior on our linearised model's weights, similarly to Fortuin et al. (2021). Specifically, we introduce dependencies between parameters belonging to the same CNN convolutional filter as

$$[\Sigma_{\theta\theta}]_{kij,k'ij'} = \sigma_d^2 \exp\left(\frac{-d(i-i', j-j')}{\ell_d}\right) \delta_{kk'}, \quad (14)$$

where  $k$  indexes the convolutional filters in the CNN,  $\delta_{kk'}$  denotes Kronecker symbol, and  $(i, j)$  index the spatial locations of specific parameters within a filter. The lengthscale  $\ell_d$  regulates the filter smoothness. Intuitively, an image generated from convolutions with smoother filters will present lower TV. Indeed, in appendix D.3 we show a bijective relationship between this quantity and the filter lengthscale. The hyperparameter  $\sigma_d^2$  determines the marginal prior variance. Both parameters are defined per architectural block  $d \in \{1, 2, \dots, D\}$  in the U-Net and we write  $\ell = [\ell_1, \ell_2, \dots, \ell_D]$  and  $\sigma^2 = [\sigma_1^2, \sigma_2^2, \dots, \sigma_D^2]$ . The chosen U-Net architecture is fully convolutional and thus eq. (14) applies to all parameters, reducing to a diagonal covariance for  $1 \times 1$  convolutions. A U-Net diagram highlighting these prior blocks is in fig. 3.

To enforce TV-smoothness, we adopt the strategy outlined in section 3. Since choosing a large  $\ell$  enforces smoothness in the output, a prior placed over the filter lengthscales  $\ell$  can act as a surrogate for the TV prior. To make this connection explicit, we construct a TV-PredCP (Nalisnick et al., 2021)

$$p(\ell) = \prod_{d=1}^D p(\ell_d) = \prod_{d=1}^D \text{Exp}(\kappa_d) \left| \frac{\partial \kappa_d}{\partial \ell_d} \right|, \quad (15)$$

$$\text{with } \kappa_d := \mathbb{E}_{\theta \sim \mathcal{N}(\hat{\theta}_d, \Sigma_{\theta_d \theta_d})} \prod_{i=1, i \neq d}^D \delta(\theta_i - \hat{\theta}_i) [\lambda \text{TV}(h(\theta))], \quad (16)$$

where  $\kappa_d$  is the expected TV of the CNN output over the prior uncertainty in the parameters of block  $d$  when all other entries of  $\theta$  are fixed to  $\hat{\theta}$ . That is, the non-smoothness introduced by a particular layer. We relate the expected TV to the filter lengthscale  $\ell_d$  by means of the change of variables formula. The independent-across-blocks of  $p(\ell)$  ensures dimensionality preservation, formally needed in the change of variables. By the triangle inequality, it can be verified that  $\sum_d \kappa_d$  is an upper bound on the expectation under the joint distribution  $\theta \sim \mathbb{E}_{\mathcal{N}(\hat{\theta}, \Sigma_{\theta\theta})} [\text{TV}(h(\theta))]$ ; see appendix D.2 for a proof.

Note that eq. (15) can be computed analytically. However, its direct computation is costly and we instead rely on numerical methods described in section 5. In fig. 2 (cf. plot 2 and plots 6-7), we show samples from  $\mathcal{N}(x; 0, \Sigma_{xx})$  where  $\ell$  is chosen using the marginal likelihood with TV-PredCP constraints (discussed subsequently in section 4.4). Incorporating the TV-PredCP leads to smoother samples with less discontinuities.

### 4.4 Type-II MAP learning of hyperparameters

The calibration of our predictive Gaussian error-bars, as described above, depends crucially on the choice of the hyperparameters of the hierarchical model given in eq. (12) (Antorán et al., 2022), that is, on the values of  $(\sigma_y^2, \sigma^2, \ell)$ . For a given lengthscale  $\ell$ , Gaussian-linear conjugacy leads to a closed form marginal likelihood objective to learn these hyperparameters. In turn, to learn  $\ell$ , we combine the aforementioned objective with the




Figure 3: A schematic of the U-Net architecture used in the numerical experiments on Walnut data (see section 6.2). For KMNIST, we use a reduced, 3-scale U-Net without group norm layers (see fig. 20). Each light-blue rectangle corresponds to a multi-channel feature map. We highlight the architectural components corresponding to each block  $1, \dots, D$  for which a separate prior is defined with red and yellow boxes.



TV-PredCP’s log-density, which acts as a regulariser. The resulting expression resembles a Type-II *maximum a posteriori* (MAP) (Rasmussen & Williams, 2005) objective

$$\log p(y|\ell; \sigma_y^2, \sigma^2) + \log p(\ell; \sigma^2) \approx -\frac{1}{2}\sigma_y^{-2}\|y - Ax(\hat{\theta})\|_2^2 - \frac{1}{2}\hat{\theta}_h^\top \Sigma_{\theta\theta}^{-1}(\ell, \sigma^2)\hat{\theta}_h - \frac{1}{2}\log |\Sigma_{yy}| - \sum_{d=1}^D \kappa_d(\ell, \sigma^2) + \log \left| \frac{\partial \kappa_d(\ell, \sigma^2)}{\partial \ell_d} \right| + B, \quad (17)$$

where  $B$  is a constant independent of the hyperparameters and the vector  $\hat{\theta}_h \in \mathbb{R}^{d_\theta}$  is the posterior mean of the linear model’s parameters. See appendix C.2 in the Appendix for the detailed derivation of the expression. The bottleneck in evaluating eq. (17) is the log-determinant of the observation covariance  $\Sigma_{yy}$ , which has a cost  $\mathcal{O}(d_y^3)$ . In the following section we describe scalable ways to approximate the log-determinant together with other expensive to compute quantities required for prediction.

## 5 Towards scalable computation

In a typical tomography setting, the dimensionality  $d_x$  of the reconstructed image  $\hat{x}$  and  $d_y$  of the observation  $y$  can be large, e.g.  $d_x > 1e5$  and  $d_y > 1e3$ . Due to the former, holding in memory the input space covariance matrices (e.g.  $\Sigma_{xx}$  and  $\Sigma_{x|y}$ ) is infeasible. The latter greatly complicates the computation of the log-determinant of the observation space covariance  $\Sigma_{yy}$  in the marginal likelihood eq. (17) (or its gradients), and its inverse in the posterior predictive distribution eq. (13), which scale as  $\mathcal{O}(d_y^3)$  and  $\mathcal{O}(d_x d_y^2)$ , respectively. To scale our approach to tomographic problems, throughout this paper, we only access Jacobian and covariance matrices through matrix–vector products, commonly known as *matvecs*. Specifically, our workhorses are products resembling  $v_x^\top \Sigma_{xx}$  and  $v_y^\top \Sigma_{yy}$  for  $v_x \in \mathbb{R}^{d_x}$  and  $v_y \in \mathbb{R}^{d_y}$ . We compute the former through successive matvecs with the components of  $\Sigma_{yy}$

$$v_y^\top \Sigma_{yy} = v_y^\top \left( AJ \Sigma_{\theta\theta} J^\top A^\top + \sigma_y^2 \mathbf{I}_{d_y} \right), \quad (18)$$

and we compute  $v_x^\top \Sigma_{xx}$  similarly. We compute Jacobian vector products  $v_\theta^\top J^\top$  for  $v_\theta \in \mathbb{R}^{d_\theta}$  using forward mode automatic differentiation (AD) and  $v_x^\top J$  using backward mode AD, both with the `functorch` library (Horace He, 2021). We compute products with  $\Sigma_{\theta\theta}$  by exploiting its block diagonal structure. All these operations can be batched using modern numerical libraries and GPUs.

### 5.1 Conjugate gradient log-determinant gradients

For the TypeII-MAP optimisation in eq. (17), we estimate the gradients of  $\log |\Sigma_{yy}|$  with respect to the parameters of interest  $\phi$  using the stochastic trace estimator (Gibbs & MacKay, 1996; Gardner et al., 2018)

$$\frac{\partial \log |\Sigma_{yy}|}{\partial \phi} = \text{Tr} \left( \Sigma_{yy}^{-1} \frac{\partial \Sigma_{yy}}{\partial \phi} \right) = \mathbb{E}_{v \sim \mathcal{N}(0, P)} \left[ v^\top \Sigma_{yy}^{-1} \frac{\partial \Sigma_{yy}}{\partial \phi} P^{-1} v \right], \quad (19)$$

where  $P$  is a preconditioner matrix. We approximately solve the linear system  $v^\top \Sigma_{yy}^{-1}$  for batches of probe vectors  $v$  using the `GPYtorch` preconditioned conjugate gradient (PCG) implementation (Dong et al., 2017).

Our preconditioner  $P$  is constructed using  $r$ -rank randomised SVD. Specifically, we approximate  $AJ \Sigma_{\theta\theta} J^\top A^\top$  as  $\tilde{U} \Lambda \tilde{U}^\top$ , using a randomised eigendecomposition algorithm (Halko et al., 2011; Martinsson & Tropp, 2020) with  $\tilde{U} \in \mathbb{R}^{d_y \times r}$  and  $r = 200 \ll d_y$ . The algorithm is described in detail in appendix F.2. Since  $P$  depends on the hyperparameters  $\phi$ , we interleave the updates of  $P$  with the optimisation of eq. (17).

### 5.2 Ancestral sampling for TV-PredCP optimisation

For large images, exact evaluation of the expected TV with eq. (16) is computationally intractable. Instead, we estimate the gradient of  $\kappa_d$  with respect to  $\phi = (\sigma^2, \ell)$  using a Monte-Carlo approximation

$$\frac{\partial \kappa_d}{\partial \phi} = \mathbb{E}_{\theta_d \sim \mathcal{N}(\hat{\theta}_d, \Sigma_{\theta_d \theta_d})} \left[ \frac{\partial \text{TV}(x)}{\partial x} J_d \frac{\partial \theta_d}{\partial \phi} \right], \quad (20)$$

where  $\frac{\partial \text{TV}(x)}{\partial x}$  is evaluated at the sample  $x = J_d \theta_d$  and  $\frac{\partial \theta_d}{\partial \phi}$  is the reparametrisation gradient for  $\theta_d$ , a prior sample of the weights of CNN block  $d$ . Since the second derivative of the TV semi-norm is almost everywhere zero, the gradient for the change of variables volume ratio is

$$\frac{\partial^2 \kappa_d}{\partial \phi^2} = \mathbb{E}_{\theta_d \sim \mathcal{N}(\hat{\theta}_d, \Sigma_{\theta_d \theta_d})} \left[ \frac{\partial \text{TV}(x)}{\partial \phi} J_d \frac{\partial^2 \theta_d}{\partial \phi^2} \right]. \quad (21)$$

### 5.3 Posterior covariance matrix estimation by sampling

The covariance matrix  $\Sigma_{x|y}$  is too large to fit into memory for high-resolution tomographic reconstructions. Instead, we follow [Wilson et al., (2021)] in drawing samples from  $\mathcal{N}(x; 0, \Sigma_{x|y})$  via Matheron’s rule

$$x_{x|y} = x + \Sigma_{xy} \Sigma_{yy}^{-1} (\epsilon - A x_0); \quad x_0 = J \theta_0; \quad \theta_0 \sim \mathcal{N}(0, \Sigma_{\theta \theta}); \quad \epsilon \sim \mathcal{N}(0, \sigma_y^2 I). \quad (22)$$

The biggest cost lies in constructing  $\Sigma_{yy}$ , which is achieved by applying eq. (18) to the standard basis vectors  $\Sigma_{yy} = [e_1, e_2, \dots, e_{d_y}]^\top \Sigma_{yy}$ . We then perform its Cholesky factorisation as an intermediate step towards matrix inversion, both relatively costly operations. Fortunately, we only have to repeat these once, after which the sampling step in eq. (22) can be evaluated cheaply. Alternatively, as in eq. (19), we can compute the solution of the linear system,  $\Sigma_{yy}^{-1} v_y$  for any  $v_y$  via PCG, without explicitly assembling (and thus storing in memory) the measurement covariance matrix, or computing its Cholesky factorisation. This approach allows us to scale the sampling operation to large measurement spaces, where the matrix  $\Sigma_{yy}$  may not fit in memory.

We now note that we only expect our predictions to be correlated for nearby pixels. Thus, we estimate cross covariances for patches of only up to  $10 \times 10$  adjacent pixels. Using larger patches yields no improvements. Specifically, we use the stabilised formulation of [Maddox et al., (2019)]:  $\hat{\Sigma}_{x|y} = \frac{1}{2k} \left[ \sum_{j=1}^k x_j^2 + x_j x_j^\top \right]$  for  $(x_j)_{j=1}^k$  samples from the posterior predictive over a patch. Note that samples from eq. (22) are zero mean.

### 5.4 Faster low-rank Jacobian matvecs

Table 1 shows that the Jacobian matvecs—implemented through forward and backward mode AD—required for sampling from the posterior predictive (that is  $2 \times v_\theta^\top J^\top$  and  $1 \times v_x^\top J$ ) take  $\approx 100\%$  of this step’s computation time (2.4 h). To accelerate sampling, we construct a low-rank approximation of the Jacobian matrix  $\tilde{J}$ , which we store in memory explicitly. We compute  $v_\theta^\top \tilde{J}^\top$  and  $v_x^\top \tilde{J}$  via matrix multiplication, as opposed to automatic differentiation. This allows for fast approximate computation of  $v_y^\top \Sigma_{yy}$  by substituting  $\tilde{J}$  into eq. (18). In turn, this brings the time needed for sampling from the posterior predictive, down from 2.4 hours to less than a minute (see table 1). We construct  $\tilde{J}$  similarly to the low-rank preconditioner  $P$  (see section 5.1 and appendix F.2). That is, following [Halko et al., (2011)], we build a structured  $r$ -rank approximation to  $J$ , by having access only to matvecs with  $J$  and  $J^\top$ .

Table 1: Wall-clock time on an A100 GPU for the different steps of our algorithm when applied to high resolution CT (details in section 6.2). Computations reported below the dotted line are in double precision. The time taken by Jacobian matvecs during sampling is given in parenthesis.

	wall-clock time
DIP optim. (after pretraining [Barbano et al., (2021b)])	<0.1 h
Hyperparam. optim. (MLL)	26.2 h
Hyperparam. optim. (TV-MAP)	35.4 h
-----	
Assemble $\Sigma_{yy}$	2.7 h
Draw 4096 posterior samples	2.4 h
- (Evaluate 4096 times $2 \times v_\theta^\top J^\top + 1 \times v_x^\top J$ )	2.4 h
Draw 4096 posterior samples ( $\tilde{J}$ & PCG)	0.3 h
- (Evaluate 4096 times $2 \times v_\theta^\top \tilde{J}^\top + 1 \times v_x^\top \tilde{J}$ )	< 0.1 min

## 6 Experiments

In this section, we experimentally evaluate: i) the properties of the models and priors discussed in sections 3 and 4, and whether they lead to accurate reconstructions and calibrated uncertainty; ii) the fidelity of the approximations described in section 5; and iii) the performance of the proposed method “linearised-DIP” (lin.-DIP) relative to the previous MC dropout (MCDO) based probabilistic formulation of DIP [Laves et al., (2020)]. We attempted to include DIP-SGLD [Cheng et al., (2019)] in our analysis, but were unable to get the method to produce competitive results on tomographic reconstruction problems. For each individual image



to be reconstructed, we employ the following linearised DIP inference procedure: i) optimise the DIP weights via eq. (5), obtaining  $\hat{x} = x(\hat{\theta})$ ; ii) optimise prior hyperparameters ( $\sigma_y^2$ ,  $\ell$ ,  $\sigma^2$ ) via eq. (17); iii) assemble and Cholesky decompose  $\Sigma_{yy}$  with eq. (18) (this step can be accelerated using approximate methods sections 5.3 and 5.4); iv) compute posterior covariance matrices either via eq. (13), or estimate them via eq. (22).

### 6.1 Small scale ablation analysis: reconstruction of KMNIST digits

Our initial analysis uses simulated CT data obtained by applying eq. (1) to 50 images from the test set of the Kuzushiji-MNIST (KMNIST) dataset:  $28 \times 28$  ( $d_x = 784$ ) grayscale images of Hiragana characters (Clanuwat et al., 2018). For each image, we choose the noise standard deviation to be either 5% or 10% of the mean of  $Ax$ , denoted as  $\eta(5\%)$  or  $\eta(10\%)$ . The forward operator  $A$  is taken to be the discrete Radon transform, assembled via ODL (Adler et al., 2017), a commonly employed software package in CT reconstruction. We use a U-Net with depth of 3 and 76905 parameters (a down-sized net compared to the one in fig. 3).

#### 6.1.1 Comparing linearised DIP with network-free priors

We first evaluate the priors described in section 3, that is, the intractable TV prior, the proposed TV-PredCP with a Matern-1/2 kernel, and a factorised Gaussian prior, by performing inference in the setting where the operator  $A$  collects 5 angles ( $d_y = 205$ ) sampled uniformly from  $0^\circ$  to  $180^\circ$  and is applied to 50 KMNIST test set images. Here, 10% noise is added. This results in a very ill-posed reconstruction problem, maximising the relevance of the prior. We select the  $\sigma_y^2$  and  $\lambda$  hyperparameters for the factorised Gaussian prior and the intractable TV prior respectively such that the posterior mean’s PSNR is maximised across a validation set of 10 images from the KMNIST training set. We keep the choice of  $\sigma_y^2$  and  $\lambda$  hyperparameters from the first two models for our experiments with the third model: Matern-1/2 with TV-PredCP prior over  $\ell$ . For all priors, we perform inference with the NUTS HMC sampler. We run 5 independent chains for each image. We burn these in for  $3 \times 10^3$  steps each and then proceed to draw  $10^4$  samples with a thinning factor of 2.

We evaluate test log-likelihood using Gaussian Kernel Density Estimation (KDE) (Silverman, 1986). The kernel bandwidth is chosen using cross-validation on 10 images from the training set. The results in table 2 show that the TV-PredCP performs best in terms of the test log-likelihood and both posterior mean and posterior mode PSNR, followed by the TV and then the factorised Gaussian. This is somewhat surprising considering that this prior was designed as an approximation to the intractable TV prior. We hypothesise that this may be due to the Matern model allowing for faster transitions in the image than the TV prior, while still capturing local correlations, as shown qualitatively in fig. 2. This property may be well-suited to the KMNIST datasets, where most pixels either present large amplitudes or are close to 0. DIP-based predictions provide 2dB higher PSNR reconstructions than the non-DIP based priors, thus linearised DIP handily obtains a better test log-likelihood than the more-traditional methods.

Table 2: Quantitative results for inference with the different priors introduced in section 3. We report both the PSNR of  $\mathbb{E}[x|y]$ , which denotes the posterior mean reconstruction, and the PSNR of  $\hat{x}$ , which denotes the posterior mode found through optimisation.

	log-likelihood	$\mathbb{E}[x y]$	$\hat{x}$
Fact. Gauss.	$0.30 \pm 0.17$	$16.15 \pm 0.38$	$14.89 \pm 0.38$
TV	$0.49 \pm 0.14$	$16.32 \pm 0.38$	$16.29 \pm 0.41$
TV-PredCP	<b><math>0.65 \pm 0.12</math></b>	<b><math>16.55 \pm 0.39</math></b>	<b><math>17.48 \pm 0.39</math></b>
lin.-DIP (MLL)	$1.63 \pm 0.08$	—	$19.46 \pm 0.52$
lin.-DIP (TV-MAP)	$1.63 \pm 0.09$	—	$19.46 \pm 0.52$

#### 6.1.2 Comparing calibration with DIP uncertainty quantification baselines

Using KMNIST, we construct test cases of different ill-posedness by simulating the observation  $y$  with four different angle sub-sampling settings for the linear operator  $A$ : 30 ( $d_y=1230$ ), 20 ( $d_y=820$ ), 10 ( $d_y=410$ ) and 5 ( $d_y=205$ ) angles are taken uniformly from the range  $0^\circ$  to  $180^\circ$ . We consider two noise configurations by adding either 5% or 10% noise to the exact data  $Ax$ . We evaluate all DIP-based methods using the same 50 randomly chosen KMNIST test set images. To ensure a best-case showing of the methods, we choose appropriate hyperparameters for each number of angles and white noise percentage setting by applying grid-search cross-validation, using 50 images from the KMNIST training dataset. Specifically, we tune the TV strength  $\lambda$  and the number of optimisation iterations for the DIP. Due to the reduced image size, we apply linearised DIP as in section 4, without approximate computations. As an ablation study, we include




Figure 4: Exemplary character recovered from  $y$  (using 5 angles and  $\eta(5\%)$ ) with lin.-DIP and DIP-MCDO along with respective uncertainty estimates. lin.-DIP provides vastly improved uncertainty calibration.

additional baselines: linearised DIP without the TV-PredCP prior over hyperparameters (labelled MLL), and DIP reconstruction with a simple Gaussian noise model consisting of the back-projected observation noise  $\mathcal{N}(x; \hat{x}, \sigma_A^2 \mathbf{I})$ , with  $\sigma_A^2 = \sigma_y^2 \text{Tr}((A^\top A)^\dagger) d_x^{-1}$  where  $\sigma_y^2=1$  (labelled  $\sigma_y^2=1$ ). Note that non-dropout methods share the same DIP parameters  $\hat{\theta}$ , and thus the same mean reconstruction. Hence, higher values in log-density indicate better uncertainty calibration, i.e., the predictive standard deviation better matches the empirical reconstruction error. DIP-MCDO does not provide an explicit likelihood function over the reconstructed image. We model its uncertainty with a Gaussian predictive distribution with covariance estimated from  $2^{14}$  samples. MNIST images are quantised to 256 bins, but our models make predictions over continuous pixel values. Thus, we simulate a de-quantisation of KMNIST images by adding a noise jitter term of variance approximately matching that of a uniform distribution over the quantisation step (Hooeboom et al., 2020).

Table 3: Mean and std-err of test log-likelihood computed over 50 KMNIST test images.

$\eta$ (5%)	#angles: 5	10	20	30	$\eta$ (10%)	#angles: 5	10	20	30
DIP ( $\sigma_y^2 = 1$ )	$0.68 \pm 0.14$	$1.57 \pm 0.02$	$1.85 \pm 0.02$	$2.02 \pm 0.02$	DIP ( $\sigma_y^2 = 1$ )	$0.27 \pm 0.17$	$1.31 \pm 0.04$	$1.62 \pm 0.03$	$1.76 \pm 0.04$
DIP-MCDO	$0.74 \pm 0.13$	$1.60 \pm 0.02$	$1.87 \pm 0.02$	$2.05 \pm 0.02$	DIP-MCDO	$0.42 \pm 0.14$	$1.39 \pm 0.04$	$1.70 \pm 0.03$	$1.85 \pm 0.04$
lin.-DIP (MLL)	$1.90 \pm 0.14$	$2.57 \pm 0.09$	$2.94 \pm 0.10$	$3.09 \pm 0.12$	lin.-DIP (MLL)	$1.63 \pm 0.08$	$2.11 \pm 0.07$	$2.43 \pm 0.07$	$2.59 \pm 0.08$
lin.-DIP (TV-MAP)	$1.88 \pm 0.15$	$2.59 \pm 0.10$	$2.96 \pm 0.10$	$3.11 \pm 0.12$	lin.-DIP (TV-MAP)	$1.63 \pm 0.09$	$2.13 \pm 0.07$	$2.45 \pm 0.08$	$2.61 \pm 0.08$

Table 4: PSNR [dB] / SSIM of the reconstruction posterior mean, averaged over 50 KMNIST test images.

$\eta$ (5%)	#angles: 5	10	20	30	$\eta$ (10%)	#angles: 5	10	20	30
DIP	<b>21.42 / 0.890</b>	27.92 / <b>0.977</b>	31.21 / <b>0.988</b>	32.93 / <b>0.991</b>	DIP	<b>19.46 / 0.846</b>	24.56 / <b>0.956</b>	27.27 / <b>0.974</b>	28.57 / <b>0.980</b>
DIP-MCDO	20.95 / 0.882	<b>28.26 / 0.977</b>	<b>31.65 / 0.986</b>	<b>33.45 / 0.990</b>	DIP-MCDO	18.91 / 0.830	<b>24.76 / 0.953</b>	<b>27.72 / 0.972</b>	<b>29.09 / 0.978</b>

Table 3 shows the test log-likelihood for all the methods and experimental settings under consideration. The peak signal-to-noise ratio (PSNR) and Structural Similarity Index (SSIM) of posterior mean reconstructions are given in table 4. All methods show similar PSNR with the standard DIP (with TV regularisation) obtaining better PSNR in the very ill-posed setting (5 angles) and MCDO obtaining marginally better reconstruction in all others. Despite this, the linearised DIP provides significantly better uncertainty calibration, outperforming all baselines in terms of test log-likelihood in all settings. Figure 4 shows an exemplary character recovered from a simulated observation  $y$  (using 20 angles and 5% noise) with both linearised DIP and DIP-MCDO along with their associated uncertainty maps and calibration plots. DIP-MCDO systematically underestimates uncertainty for pixels on which the error is large, explaining its poor test log-likelihood. The pixel-wise standard deviation provided by linearised DIP (TV-MAP) better correlates with the reconstruction error.

### 6.1.3 Evaluating the fidelity of sample-based predictive covariance matrix estimation

We evaluate the accuracy of the sampling, conjugate gradient and low rank approximations to the predictive covariance  $\Sigma_{x|y}$  discussed in section 5. As a reference, we compute the exact predictive covariance with eq. (13), which is tractable for KMNIST. We do not use patch-based approximations or stabilised covariance

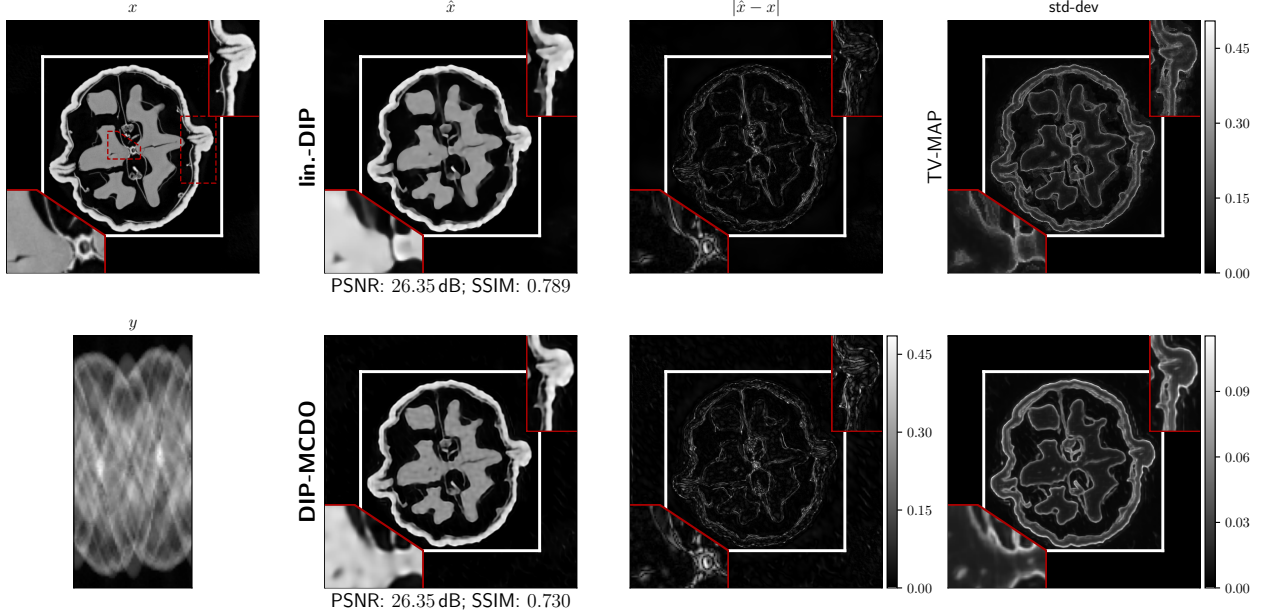


Figure 5: Reconstruction of a  $501 \times 501 \text{ px}^2$  slice of a scanned Walnut using lin.-DIP and DIP-MCDO along with their respective uncertainty estimates. The zoomed regions (outlined in red) are given in top-left.

estimators for this experiment. Table 5 shows that estimating the covariance matrix using samples provides no decrease in performance. Using a low-rank approximation to the Jacobian matrix together with computing linear solves with PCG loses at most 0.32 nats in test log-likelihood with respect to the exact computation, but results in almost an order of magnitude speedup at prediction time.

Table 5: Evaluation of our approximate covariance estimation methods in terms of test log-likelihood over 10 KMNIST test images considering the 20 angle ( $d_y = 820$ ) setting and using lin.-DIP (MLL).

	exact	sampler	sampler	sampler cov.
$\eta$ (%)	cov. eq. [13]	cov. eq. [22]	cov. ( $\tilde{J}$ ) eq. [22]	( $\tilde{J}$ & PCG) eq. [22]
5	$2.80 \pm 0.06$	$2.80 \pm 0.06$	$2.68 \pm 0.09$	$2.62 \pm 0.09$
10	$2.26 \pm 0.06$	$2.26 \pm 0.06$	$2.21 \pm 0.06$	$2.22 \pm 0.06$

Table 6: Test log-likelihood, PSNR and structural similarity (SSIM) on the Walnut. We compare all lin.-DIP variants with DIP-MCDO.

	$1 \times 1$	$2 \times 2$	$10 \times 10$	PSNR [dB]	SSIM
DIP-MCDO	0.03	1.68	2.47	23.49	0.730
lin.-DIP (MLL)	2.09	2.25	2.43	<b>26.35</b>	<b>0.789</b>
lin.-DIP (MLL, $\tilde{J}$ & PCG)	1.88	2.05	2.24	—	—
lin.-DIP (TV-MAP)	<b>2.21</b>	<b>2.40</b>	<b>2.60</b>	—	—
lin.-DIP (TV-MAP, $\tilde{J}$ & PCG)	<b>2.24</b>	<b>2.46</b>	<b>2.65</b>	—	—

## 6.2 Linearised DIP for high-resolution CT

We now demonstrate our approach on real-measured cone-beam  $\mu$ CT data obtained by scanning a walnut (Der Sarkissian et al., 2019). We reconstruct a  $501 \times 501 \text{ px}^2$  slice ( $d_x = 251 \cdot 10^3$ ) using a sparse subset of measurements taken from 60 angles and 128 detector rows ( $d_y = 7680$ ). We use the U-Net shown in fig. 3 which contains approximately 3 million parameters. Here,  $\Sigma_{xx}$  is too large to store in memory and  $\Sigma_{yy}$  is too expensive to assemble repeatedly, leading us to use the full suite of approximate computations introduced in section 5. Since the Walnut data is not quantised, jitter correction is not needed.

As shown in appendix E, during MLL and Type-II MAP optimisation, many layers’ prior variance goes to  $\sigma_d^2 \approx 0$ . This phenomenon is known as “automatic relevance determination” (Mackay, 1996; Tipping,




Figure 6: Optimisation trajectories for hyperparameters of the U-net’s first and last  $3 \times 3$  convolutions ( $\ell_1, \sigma_{\ell_1}^2, \ell_{11}, \sigma_{\ell_{11}}^2$ ), last  $1 \times 1$  convolution ( $\sigma_{1 \times 1,3}^2$ ) and noise variance  $\sigma_y^2$  for the Walnut data.




Figure 7: The comparison of uncertainty calibration: the pixel-wise reconstruction absolute error  $|\hat{x} - x|$  overlaps with the uncertainties provided by the lin.-DIP. DIP-MCDO, instead, severely underestimates uncertainty. The scale of the pixel-wise standard deviation (std-dev) obtained including the TV-PredCP matches the absolute error more closely than when the hyperparameters are optimised without. Using  $\tilde{J}$  & PCG results in overestimating uncertainty in the tails. LL stands for test log-likelihood.

[2001]), and simplifies our linearised network, preventing uncertainty overestimation. We did not observe this effect when working with KMNIST images and smaller networks. We display the MLL and MAP optimisation profiles for the active layers (i.e. layers with high  $\sigma_d^2$ ) in fig. 6. Type-II MAP hyperparameters optimisation drives  $\sigma^2$  to smaller values, compared to MLL. This restricts the linearised DIP prior, and thus the induced posterior, to functions that are smooth in a TV sense, leading to smaller error-bars c.f. fig. 7. As the optimisation of eq. (17) progresses,  $\ell_1, \ell_{11}$  fall into basins of new minima corresponding to larger lengthscales. This results in more correlated dimensions in the prior, further simplifying the model.

Density estimation operations, described in section 5.3, are conducted in double precision (64 bit floating point) as we found single precision led to numerical instability in the assembly of  $\Sigma_{yy}$ , and also in the estimation of off-diagonal covariance terms for larger patches. In table 6, we report test log-likelihood computed using a Gaussian predictive distribution with covariance patches of sizes  $1 \times 1$ ,  $2 \times 2$  and  $10 \times 10$  pixels. Mean reconstruction metrics are also reported. Figure 5 displays reconstructed images, uncertainty maps and calibration plots. In this more challenging tomographic reconstruction task, DIP-MCDO performs poorly relative to the standard regularised DIP formulation eq. (5) in terms of reconstruction PSNR. DIP-MCDO underestimates uncertainty, and its uncertainty map is blurred across large sections of the image, placing large uncertainty in well-reconstructed regions and vice-versa. In contrast, the uncertainty map provided by linearised DIP is fine-grained, concentrating on regions of increased reconstruction error. Quantitatively, linearised DIP provides over 2.06 nats per pixel improvement in terms of test log-likelihood and more calibrated uncertainty estimates, as reflected in the Q-Q plot in fig. 7. Furthermore, the use of our TV-PredCP based prior for MAP optimisation yields a 0.12 nat per pixel improvement over the MLL approach. Interestingly, using low-rank Jacobians and PCG for sampling provides a small performance boost when using the TV-PredCP prior. Figure 7 reveals that these approximations result in uncertainty overestimation (a known issue (Antoran et al., 2023)) which is compensated by the more restrictive TV-PredCP prior.

## 7 Conclusion

We have proposed a probabilistic formulation of the deep image prior (DIP) that utilises a linearisation of the DIP parameters around the mode of the loss and a Gaussian-linear hierarchical prior on the network parameters mimicking the total variation prior (constructed via the predictive complexity prior framework). The approach yields well-calibrated uncertainty estimates on tomographic reconstruction tasks based on simulated observations as well as on real-measured  $\mu$ CT data. The empirical results suggest that the both the DIP reparametrisation and the TV regulariser provide good inductive biases for both high-quality reconstructions and well-calibrated uncertainty estimates. The proposed method is shown to provide by far more calibrated uncertainty estimates than existing MC dropout approaches to uncertainty estimation with DIP.

## References

- J Adler, H Kohr, and O Oktem. Operator discretization library (ODL). *Software available from* <https://github.com/odlgroup/odl>, 2017.
- Javier Antorán, James Allingham, and José Miguel Hernández-Lobato. Depth uncertainty in neural networks. *Advances in Neural Information Processing Systems*, 33:10620–10634, 2020.
- Javier Antorán, Umang Bhatt, Tameem Adel, Adrian Weller, and José Miguel Hernández-Lobato. Getting a {clue}: A method for explaining uncertainty estimates. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=XSLF1XFq5h>.
- Javier Antorán, James Allingham, David Janz, Erik Daxberger, Eric Nalisnick, and José Miguel Hernández-Lobato. Linearised Laplace inference in networks with normalisation layers and the neural g-prior. Fourth Symposium on Advances in Approximate Bayesian Inference, AABI 2022, 2022.
- Javier Antorán, David Janz, James Urquhart Allingham, Erik A. Daxberger, Riccardo Barbano, Eric T. Nalisnick, and José Miguel Hernández-Lobato. Adapting the linearised Laplace model evidence for modern deep learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 796–821. PMLR, 2022. URL <https://proceedings.mlr.press/v162/antoran22a.html>.
- Javier Antoran, Shreyas Padhy, Riccardo Barbano, Eric Nalisnick, David Janz, and José Miguel Hernández-Lobato. Sampling-based inference for large linear models, with application to linearised laplace. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=aoDyX6vSqsdl>.
- 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. *Proc. Nat. Acad. Sci.*, 117(48):30088–30095, 2020.
- Simon Arridge, Peter Maaß, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta Numer.*, 28:1–174, 2019.
- Arsenii Ashukha, Alexander Lyzhov, Dmitry Molchanov, and Dmitry Vetrov. Pitfalls of in-domain uncertainty estimation and ensembling in deep learning. Preprint arXiv:2002.06470, 2020.
- Daniel Otero Bager, Johannes Leuschner, and Maximilian Schmidt. Computed tomography reconstruction using deep image prior and learned reconstruction methods. *Inverse Problems*, 36(9):094004, 2020.
- Riccardo Barbano, Javier Antorán, José Miguel Hernández-Lobato, and Bangti Jin. A probabilistic deep image prior over image space. In *Fourth Symposium on Advances in Approximate Bayesian Inference*.
- Riccardo Barbano, Simon Arridge, Bangti Jin, and Ryutaro Tanno. Uncertainty quantification for medical image synthesis. In *Biomedical Image Synthesis and Simulation: Methods and Applications*, pp. 601–641. Elsevier, 2021a.
- Riccardo Barbano, Johannes Leuschner, Maximilian Schmidt, Alexander Denker, Andreas Hauptmann, Peter Maaß, and Bangti Jin. An educated warm start for deep image prior-based micro ct reconstruction. Preprint, arXiv:2111.11926, 2021b.
- Riccardo Barbano, Johannes Leuschner, Javier Antorán, Bangti Jin, and José Miguel Hernández-Lobato. Bayesian experimental design for computed tomography with the linearised deep image prior. In *ICML2022 Workshop on Adaptive Experimental Design and Active Learning in the Real World*, 2022.
- Riccardo Barbano, Javier Antorán, Johannes Leuschner, José Miguel Hernández-Lobato, Željko Kereta, and Bangti Jin. Fast and painless image reconstruction in deep image prior subspaces, 2023.



- Semih Barutcu, Doaa Gürsoy, and Aggelos K. Katsaggelos. Compressive ptychography using deep image and generative priors. Preprint, arXiv:2205.02397, 2022.
- Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q. Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, Lama Nachman, Rumi Chunara, Madhulika Srikumar, Adrian Weller, and Alice Xiang. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 401–413, New York, NY, USA, 2021. Association for Computing Machinery. URL <https://doi.org/10.1145/3461702.3462571>.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. In *Proceedings of the 32nd International Conference on Machine Learning, PMLR 37*, pp. 1613–1622, 2015.
- Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed sensing using generative models. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 537–546, 2017.
- Antonin Chambolle, Vicent Caselles, Daniel Cremers, Matteo Novaga, and Thomas Pock. An introduction to total variation for image analysis. In *Theoretical foundations and numerical methods for sparse recovery*, pp. 263–340. de Gruyter, 2010.
- Zezhou Cheng, Matheus Gadelha, Subhransu Maji, and Daniel Sheldon. A Bayesian perspective on the deep image prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5443–5451. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00559.
- Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical Japanese literature. In *32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Workshop on Machine Learning for Creativity and Design*, 2018.
- Jianan Cui, Kuang Gong, Ning Guo, Chenxi Wu, Kyungsang Kim, Huafeng Liu, and Quanzheng Li. Populational and individual information based PET image denoising using conditional unsupervised learning. *Phys. Med. & Biol.*, 66(15):155001, 2021.
- Mohammad Zalbagi Darestani and Reinhard Heckel. Accelerated MRI with un-trained neural networks. *IEEE Trans. Comput. Imag.*, 7:724–733, 2021.
- Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux - effortless Bayesian deep learning. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34*, pp. 20089–20103, 2021a. URL <https://proceedings.neurips.cc/paper/2021/hash/a7c9585703d275249f30a088cebba0ad-Abstract.html>.
- Erik Daxberger, Eric Nalisnick, James U Allingham, Javier Antorán, and Jose Miguel Hernandez-Lobato. Bayesian deep learning via subnetwork inference. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, pp. 2510–2521, 2021b. URL <https://proceedings.mlr.press/v139/daxberger21a.html>.
- Henri Der Sarkissian, Felix Lucka, Maureen van Eijnatten, Giulia Colacicco, Sophia Bethany Coban, and K. Joost Batenburg. Cone-Beam X-Ray CT Data Collection Designed for Machine Learning: Samples 1-8, 2019. URL <https://doi.org/10.5281/zenodo.2686726>. *Zenodo*.
- Kun Dong, David Eriksson, Hannes Nickisch, David Bindel, and Andrew Gordon Wilson. Scalable log determinants for Gaussian process kernel learning. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6327–6337, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/976abf49974d4686f87192efa0513ae0-Abstract.html>.



- Heinz W. Engl, Martin Hanke, and Andreas Neubauer. *Regularization of Inverse Problems*. Kluwer, Dordrecht, 1996. ISBN 0-7923-4157-0.
- Andrew Y. K. Foong, David R. Burt, Yingzhen Li, and Richard E. Turner. On the expressiveness of approximate inference in Bayesian neural networks. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/b6dfd41875bc090bd31d0b1740eb5b1b-Abstract.html>.
- Vincent Fortuin, Adrià Garriga-Alonso, Florian Wenzel, Gunnar Ratsch, Richard E Turner, Mark van der Wilk, and Laurence Aitchison. Bayesian neural network priors revisited. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2021. URL <https://openreview.net/forum?id=xaqKWHcoOGP>.
- Jacob R. Gardner, Geoff Pleiss, Kilian Q. Weinberger, David Bindel, and Andrew Gordon Wilson. GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 7587–7597, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/27e8e17134dd7083b050476733207ea1-Abstract.html>.
- Adrià Garriga-Alonso, Laurence Aitchison, and Carl Edward Rasmussen. Deep convolutional networks as shallow Gaussian processes. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bklfsi0cKm>.
- M. N. Gibbs and D. J. C. MacKay. Efficient implementation of Gaussian processes for interpolation. <http://www.inference.phy.cam.ac.uk/mackay/abstracts/gpros.html>, 1996.
- Kuang Gong, Ciprian Catana, Jinyi Qi, and Quanzheng Li. PET image reconstruction using deep image prior. *IEEE Trans. Med. Imag.*, 38(7):1655–1665, 2019.
- Peter Guttorp and Tilmann Gneiting. On the Whittle-Matérn correlation family. NRCSE Technical Report No. 80, University of Washington, 01 2005.
- Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011. doi: 10.1137/090771806. URL <https://doi.org/10.1137/090771806>.
- Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. In *ICLR*, 2019. URL <https://openreview.net/pdf?id=rylV-2C9KQ>.
- Tapio Helin, Nuutti Hyvönen, and Juha-Pekka Puska. Edge-promoting adaptive Bayesian experimental design for x-ray imaging. *SIAM J. Sci. Comput.*, 44(3):B506–B530, 2022. doi: 10.1137/21m1409330. URL <https://doi.org/10.1137/21m1409330>.
- Emiel Hoogeboom, Taco S. Cohen, and Jakub M. Tomczak. Learning discrete distributions by dequantization. Preprint, arXiv:2001.11235, 2020.
- Richard Zou Horace He. functorch: JAX-like composable function transforms for PyTorch. <https://github.com/pytorch/functorch>, 2021.
- Jiri Hron, Alex Matthews, and Zoubin Ghahramani. Variational Bayesian dropout: pitfalls and fixes. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2019–2028. PMLR, 10–15 Jul 2018.
- Alexander Immer, Matthias Bauer, Vincent Fortuin, Gunnar Rätsch, and Mohammad Emtiyaz Khan. Scalable marginal likelihood estimation for model selection in deep learning. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4563–4573. PMLR, 2021a. URL <http://proceedings.mlr.press/v139/immer21a.html>.

- Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of Bayesian neural nets via local linearization. In Arindam Banerjee and Kenji Fukumizu (eds.), *The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 703–711. PMLR, 2021b. URL <http://proceedings.mlr.press/v130/immer21a.html>
- Kazufumi Ito and Bangti Jin. *Inverse Problems: Tikhonov Theory and Algorithms*, volume 22. World Scientific, 2014.
- Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 8580–8589, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/5a4be1fa34e62bb8a6ec6b91d2462f5a-Abstract.html>
- Jari Kaipio and Erkki Somersalo. *Statistical and Computational Inverse Problems*. Springer-Verlag, New York, 2005. ISBN 0-387-22073-9.
- Mohammad Emtiyaz Khan, Alexander Immer, Ehsan Abedi, and Maciej Korzepa. Approximate inference turns deep networks into Gaussian processes. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 3088–3098, 2019.
- Tobias Knopp and Micro Grosser. Warmstart approach for accelerating deep image prior reconstruction in dynamic tomography. *Proceedings of Machine Learning Research, Medical Imaging with Deep Learning 2022*, 13 pp., 2022.
- Benjamin Kompa, Jasper Snoek, and Andrew L Beam. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digital Medicine*, 4(1):1–6, 2021.
- Max-Heinrich Laves, Malte Tölle, and Tobias Ortmaier. Uncertainty estimation in medical image denoising with Bayesian deep image prior. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis*, pp. 81–96. 2020.
- Jaehoon Lee, Lechao Xiao, Samuel S. Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8570–8581, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/0d1a9651497a38d8b1c3871c84528bd4-Abstract.html>
- F. C. Leone, L. S. Nelson, and R. B. Nottingham. The folded normal distribution. *Technometrics*, 3:543–550, 1961. ISSN 0040-1706. doi: 10.2307/1266560. URL <https://doi.org/10.2307/1266560>.
- Jiaming Liu, Yu Sun, Xiaojian Xu, and Ulugbek S Kamilov. Image restoration using total variation regularized deep image prior. In *ICASSP 2019*, 2019. doi: 10.1109/ICASSP.2019.8682856.
- David J. C. Mackay. Bayesian non-linear modeling for prediction competition. In *Maximum Entropy and Bayesian Methods*, pp. 221–234, 1996.
- David John Cameron Mackay. *Bayesian Methods for Adaptive Models*. PhD thesis, California, USA, 1992.
- Wesley Maddox, Shuai Tang, Pablo Garcia Moreno, Andrew Gordon Wilson, and Andreas C. Damianou. Fast adaptation with linearized neural networks. In Arindam Banerjee and Kenji Fukumizu (eds.), *The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 2737–2745. PMLR, 2021. URL <http://proceedings.mlr.press/v130/maddox21a.html>
- Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for Bayesian uncertainty in deep learning. *Advances in Neural Information Processing Systems*, 32, 2019.

- Per-Gunnar Martinsson and Joel A Tropp. Randomized numerical linear algebra: Foundations and algorithms. *Acta Numer.*, 29:403–572, 2020.
- Kenneth O. McGraw and S. P. Wong. The descriptive use of absolute differences between pairs of scores with a common mean and variance. *J. Educat. Stat.*, 19(2):103–110, 1994.
- Eric T. Nalisnick, Jonathan Gordon, and José Miguel Hernández-Lobato. Predictive complexity priors. In Arindam Banerjee and Kenji Fukumizu (eds.), *The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 694–702. PMLR, 2021. URL <http://proceedings.mlr.press/v130/nalisnick21a.html>
- Roman Novak, Lechao Xiao, Yasaman Bahri, Jaehoon Lee, Greg Yang, Jiri Hron, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian deep convolutional networks with many channels are Gaussian processes. In *7th International Conference on Learning Representations*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Blg30j0qF7>.
- Roman Novak, Lechao Xiao, Jiri Hron, Jaehoon Lee, Alexander A. Alemi, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. Neural tangents: Fast and easy infinite neural networks in Python. In *8th International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=Sk1D9yrFPS>.
- Gregory Ongie, Ajil Jalal, Richard G Baraniuk, Christopher A Metzler, Alexandros G Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE J. Sel. Areas Inform. Theory*, pp. 39–56, 2020.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2005. ISBN 026218253X.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241, 2015.
- Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, 1992.
- Walter Rudin. *Fourier Analysis on Groups*. John-Wiley, New York-London, 1990.
- Matthias W. Seeger and Hannes Nickisch. Large scale Bayesian inference and experimental design for sparse linear models. *SIAM J. Imaging Sci.*, 4(1):166–199, 2011. doi: 10.1137/090758775. URL <https://doi.org/10.1137/090758775>.
- Zenglin Shi, Pascal Mettes, Subhansu Maji, and Cees G. M. Snoek. On measuring and controlling the spectral bias of the deep image prior. *Int. J. Comput. Vis.*, 130(4):885–908, 2022. doi: 10.1007/s11263-021-01572-7. URL <https://doi.org/10.1007/s11263-021-01572-7>.
- Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, London, 1986. ISBN 0-412-24620-1.
- Jasper Snoek, Yaniv Ovadia, Emily Fertig, Balaji Lakshminarayanan, Sebastian Nowozin, D. Sculley, Joshua V. Dillon, Jie Ren, and Zachary Nado. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 13969–13980, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/8558cb408c1d76621371888657d2eb1d-Abstract.html>
- Andrew M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numer.*, 19:451–559, 2010. ISSN 0962-4929. doi: 10.1017/S0962492910000061. URL <https://doi.org/10.1017/S0962492910000061>.
- Andrey N. Tikhonov and Vasilii Y. Arsenin. *Solutions of Ill-posed Problems*. John Wiley & Sons, New York-Toronto, Ont.-London, 1977.

- Michael E Tipping. Sparse Bayesian learning and the relevance vector machine. *J. Mach. Learn. Res.*, pp. 211–244, 2001.
- Malte Tölle, Max-Heinrich Laves, and Alexander Schlaefer. A mean-field variational inference approach to deep image prior for inverse problems in medical imaging. In Mattias P. Heinrich, Qi Dou, Marleen de Bruijne, Jan Lellmann, Alexander Schlaefer, and Floris Ernst (eds.), *Medical Imaging with Deep Learning, 7-9 July 2021, Lübeck, Germany*, volume 143 of *Proceedings of Machine Learning Research*, pp. 745–760. PMLR, 2021. URL <https://proceedings.mlr.press/v143/tolle21a.html>.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9446–9454, 2018.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *Int. J. Comput. Vis.*, 128(7): 1867–1888, 2020. ISSN 1573-1405. doi: 10.1007/s11263-020-01303-4.
- Wim van Aarle, Willem Jan Palenstijn, Jan De Beenhouwer, Thomas Altantzis, Sara Bals, K. Joost Batenburg, and Jan Sijbers. The ASTRA Toolbox: A platform for advanced algorithm development in electron tomography. *Ultramicroscopy*, 157:35–47, 2015. doi: <https://doi.org/10.1016/j.ultramic.2015.05.002>.
- Francisca Vasconcelos, Bobby He, Nalini Singh, and Yee Whye Teh. UncertaINR: Uncertainty quantification of end-to-end implicit neural representations for computed tomography. Preprint, arXiv:2202.10847, 2022.
- Ge Wang, Jong Chul Ye, and Bruno De Man. Deep learning for tomographic image reconstruction. *Nature Mach. Intell.*, 2(12):737–748, 2020.
- Hengkang Wang, Taihui Li, Zhong Zhuang, Tiancong Chen, Hengyue Liang, and Ju Sun. Early stopping for deep image prior. *CoRR*, abs/2112.06074, 2021. URL <https://arxiv.org/abs/2112.06074>.
- Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In Lise Getoor and Tobias Scheffer (eds.), *Proceedings of the 28th International Conference on Machine Learning*, pp. 681–688. Omnipress, 2011. URL [https://icml.cc/2011/papers/398\\_icmlpaper.pdf](https://icml.cc/2011/papers/398_icmlpaper.pdf).
- James T. Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Pathwise conditioning of Gaussian processes. *J. Mach. Learn. Res.*, 22:105:1–105:47, 2021. URL <http://jmlr.org/papers/v22/20-1260.html>.