

BayesCap: Bayesian Identity Cap for Calibrated Uncertainty in Frozen Neural Networks

Uddeshya Upadhyay^{*1}, Shyamgopal Karthik^{*1}, Yanbei Chen¹,
Massimiliano Mancini¹, and Zeynep Akata^{1,2}

¹University of Tübingen ²Max Planck Institute for Intelligent Systems

Abstract. High-quality calibrated uncertainty estimates are crucial for numerous real-world applications, especially for deep learning-based deployed ML systems. While Bayesian deep learning techniques allow uncertainty estimation, training them with large-scale datasets is an expensive process that does not always yield models competitive with non-Bayesian counterparts. Moreover, many of the high-performing deep learning models that are already trained and deployed are non-Bayesian in nature and do not provide uncertainty estimates. To address these issues, we propose **BayesCap** that learns a Bayesian identity mapping for the frozen model, allowing uncertainty estimation. **BayesCap** is a memory-efficient method that can be trained on a small fraction of the original dataset, enhancing pretrained non-Bayesian computer vision models by providing calibrated uncertainty estimates for the predictions without (i) hampering the performance of the model and (ii) the need for expensive retraining the model from scratch. The proposed method is agnostic to various architectures and tasks. We show the efficacy of our method on a wide variety of tasks with a diverse set of architectures, including image super-resolution, deblurring, inpainting, and crucial application such as medical image translation. Moreover, we apply the derived uncertainty estimates to detect out-of-distribution samples in critical scenarios like depth estimation in autonomous driving. Code is available at <https://github.com/ExplainableML/BayesCap>.

Keywords: Uncertainty Estimation, Calibration, Image Translation

1 Introduction

Image enhancement and translation tasks like super-resolution [36], deblurring [30,31], inpainting [74], colorization [78,25], denoising [58,48], medical image synthesis [81,1,11,60,59,63,57], monocular depth estimation in autonomous driving [16,21], etc., have been effectively tackled using deep learning methods generating high-fidelity outputs. But, the respective state-of-the-art models usually learn a *deterministic* one-to-one mapping between the input and the output, without modeling the uncertainty in the prediction. For instance, a depth estimation model predicts a depth map from the input RGB image (Figure 1-(Left)), without providing uncertainty. In contrast, learning a probabilistic mapping between the input and the output yields the underlying distribution and provides

^{*}Equal contribution




Fig. 1: Computer vision models for image enhancements and translations deterministically map input to output without producing the uncertainty in the latter (example on the right shows depth estimation using MonoDepth2 [21]). BayesCap approximates the underlying distribution and adds uncertainty to the predictions of pretrained models efficiently, details in Section 3.3.

uncertainty estimates for the predictions. This is a vital feature in safety-critical applications such as autonomous driving and medical imaging. For instance, well-calibrated uncertainty estimates can be used to trigger human/expert intervention in highly uncertain predictions, consequently preventing fatal automated decision making [10,54,53]. The conventional approach for obtaining uncertainty estimates is to train Bayesian models *from scratch*. However, Bayesian deep learning techniques are difficult to train and are not scalable to large volumes of high-dimensional data [12]. Moreover, they cannot be easily integrated with sophisticated deterministic architectures and training schemes tailored for specific tasks, vital to achieving state-of-the-art in vision applications [14,20].

To address the above challenges, we enhance the predictions of pretrained state-of-the-art non-Bayesian deterministic deep models with uncertainty estimation while preserving their strong model performances. There is limited literature tackling the similar problem [66,3,12] but these methods do not yield well-calibrated uncertainty estimates or do not scale to high-dimensional cases such as image synthesis, translation, and enhancement.

In this work, we propose **BayesCap**, shown in Figure 1-(Right), an architecture agnostic, plug-and-play method to generate uncertainty estimates for pretrained models. The key idea is to train a Bayesian autoencoder over the output images of the pretrained network, approximating the underlying distribution for the output images. Due to its Bayesian design, in addition to reconstructing the input, **BayesCap** also estimates the parameters of the underlying distribution, allowing us to compute the uncertainties. **BayesCap** is highly data-efficient and can be trained on a small fraction of the original dataset. For instance, **BayesCap** is 3-5 \times faster to train as compared to a Bayesian model from scratch, while still achieving uncertainty estimates that are better calibrated than the baselines.

To summarize, we make the following contributions. (1) We propose **BayesCap**, a simple method for generating post-hoc uncertainty estimates, by learning a Bayesian identity mapping, over the outputs of image synthesis/translation tasks with deterministic pretrained models. (2) **BayesCap** leads to calibrated uncertainties while retaining the performance of the underlying state-of-the-art pretrained network on a variety of tasks including super-resolution, deblurring, inpainting, and medical imaging. (3) We also show that quantifying uncertainty us-

ing BayesCap can help in downstream tasks such as Out-of-Distribution (OOD) detection in critical applications like autonomous driving.

2 Related Works

Image Enhancement and Translations. Advances in computer vision led to tackle challenging problems such as super-resolution [13,36], denoising [58,48], deblurring [44,30,31], inpainting [46,74], depth estimation [16,21] among others. Such problems are tackled using a diverse set of architectures and learning schemes. For instance, the popular method for super-resolution involves training a conditional *generative adversarial networks* (GANs), where the generator is conditioned with a low-resolution image and employs a pretrained VGG network [55] to enforce the content loss in the feature space along with the adversarial term from the discriminator [36]. Differently, for the inpainting task, [74] uses a conditional GAN with contextual attention and trains the network using spatially discounted reconstruction loss. In the case of monocular depth estimation, recent works exploit the left-right consistency as a cue to train the model in an unsupervised fashion [20]. While these methods are highly diverse in their architectures, training schemes, supervisory signals, etc., they typically focus on providing a deterministic one-to-one mapping which may not be ideal in many critical scenarios such as autonomous driving [42] and medical imaging [62,63,61]. Our BayesCap preserves the high-fidelity outputs provided by such deterministic pretrained models while approximating the underlying distribution of the output of such models, allowing uncertainty estimation.

Uncertainty Estimation. Bayesian deep learning models are capable of estimating the uncertainties in their prediction [32,27]. Uncertainties can be broadly divided into two categories; (1) Epistemic uncertainty which is the uncertainty due to the model parameters [22,6,12,70,9,32,17]. (2) Aleatoric uncertainty which is the underlying uncertainty in the measurement itself, often estimated by approximating the per-pixel residuals between the predictions and the ground-truth using a *heteroscedastic* distribution whose parameters are predicted as the output of the network which is trained *from scratch* to maximize the likelihood of the system [27,38,4,66,67,33]. While epistemic uncertainty is important in low-data regimes as parameter estimation becomes noisy, however, this is often not the case in computer vision settings with large scale datasets where aleatoric uncertainty is the critical quantity [27]. However, it is expensive to train these models and they often perform worse than their deterministic counterparts [45,50,12]. Unlike these works, BayesCap is a fast and efficient method to estimate uncertainty over the predictions of a pretrained deterministic model.

Post-hoc Uncertainty Estimation. While this has not been widely explored, some recent works [12,15] have tried to use the Laplace approximation for this purpose. However, these methods compute the Hessian which is not feasible for high-dimensional modern problems in computer vision [78,46,16,30,48]. Another line of work to tackle this problem is test-time data augmentation [66,3] that perturbs the inputs to obtain multiple outputs leading to uncertainties.

However, these estimates are often poorly calibrated [18]. It is of paramount importance that the uncertainty estimates are well calibrated [29,23,35,33,47,77]. In many high-dimensional computer vision problems the per-pixel output is often a continuous value [46,30,36,78], i.e., the problem is regression in nature. Recent works focused on *Uncertainty Calibration Error* that generalizes to high dimensional regression [29,37,35,33]. Unlike prior works [3,12,66,15], BayesCap scales to high-dimensional tasks, providing well-calibrated uncertainties.

3 Methodology: BayesCap - Bayesian Identity Cap

We first describe the problem formulation in Section 3.1, and preliminaries on uncertainty estimation in Section 3.2. In Section 3.3, we describe construction of BayesCap that models a probabilistic identity function capable of estimating the high-dimensional complex distribution from the frozen deterministic model, estimating calibrated uncertainty for the predictions.

3.1 Problem formulation

Let $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ be the training set with pairs from domain \mathbf{X} and \mathbf{Y} (i.e., $\mathbf{x}_i \in \mathbf{X}, \mathbf{y}_i \in \mathbf{Y}, \forall i$), where \mathbf{X}, \mathbf{Y} lies in \mathbb{R}^m and \mathbb{R}^n , respectively. While our proposed solution is valid for data of arbitrary dimension, we present the formulation for images with applications for image enhancement and translation tasks, such as super-resolution, inpainting, etc. Therefore, $(\mathbf{x}_i, \mathbf{y}_i)$ represents a pair of images, where \mathbf{x}_i refers to the input and \mathbf{y}_i denotes the transformed/enhanced output. For instance, in super-resolution \mathbf{x}_i is a low-resolution image and \mathbf{y}_i its high-resolution version. Let $\Psi(\cdot; \theta) : \mathbb{R}^m \rightarrow \mathbb{R}^n$ represent a Deep Neural Network parametrized by θ that maps images from the set \mathbf{X} to the set \mathbf{Y} , e.g. from corrupted to the non-corrupted/enhanced output images.

We consider a real-world scenario, where $\Psi(\cdot; \theta)$ has already been trained using the dataset \mathcal{D} and it is in a *frozen state* with parameters set to the learned optimal parameters θ^* . In this state, given an input \mathbf{x} , the model returns a point estimate of the output, i.e., $\hat{\mathbf{y}} = \Psi(\mathbf{x}; \theta^*)$. However, point estimates do not capture the distributions of the output ($\mathcal{P}_{\mathbf{Y}|\mathbf{X}}$) and thus the uncertainty in the prediction that is crucial in many real-world applications [27]. Therefore, we propose to estimate $\mathcal{P}_{\mathbf{Y}|\mathbf{X}}$ for the pretrained model in a fast and cheap manner, quantifying the uncertainties of the output without re-training the model itself.

3.2 Preliminaries: Uncertainty Estimation

To understand the functioning of our BayesCap that produces uncertainty estimates for the *frozen or pretrained* neural networks, we first consider a model trained from scratch to address the target task and estimate uncertainty. Let us denote this model by $\Psi_s(\cdot; \zeta) : \mathbb{R}^m \rightarrow \mathbb{R}^n$, with a set of trainable parameters given by ζ . To capture the *irreducible* (i.e., aleatoric) uncertainty in the output distribution $\mathcal{P}_{\mathbf{Y}|\mathbf{X}}$, the model must estimate the parameters of the distribution. These are then used to maximize the likelihood function. That is,

for an input \mathbf{x}_i , the model produces a set of parameters representing the output given by, $\{\hat{\mathbf{y}}_i, \hat{\nu}_i \dots \hat{\rho}_i\} := \Psi_s(\mathbf{x}_i; \zeta)$, that characterizes the distribution $\mathcal{P}_{Y|X}(\mathbf{y}; \{\hat{\mathbf{y}}_i, \hat{\nu}_i \dots \hat{\rho}_i\})$, such that $\mathbf{y}_i \sim \mathcal{P}_{Y|X}(\mathbf{y}; \{\hat{\mathbf{y}}_i, \hat{\nu}_i \dots \hat{\rho}_i\})$. The likelihood $\mathcal{L}(\zeta; \mathcal{D}) := \prod_{i=1}^N \mathcal{P}_{Y|X}(\mathbf{y}_i; \{\hat{\mathbf{y}}_i, \hat{\nu}_i \dots \hat{\rho}_i\})$ is then maximized in order to estimate the optimal parameters of the network. Moreover, the distribution $\mathcal{P}_{Y|X}$ is often chosen such that uncertainty can be estimated using a closed form solution \mathcal{F} depending on the estimated parameters of the neural network, i.e.,

$$\{\hat{\mathbf{y}}_i, \hat{\nu}_i \dots \hat{\rho}_i\} := \Psi_s(\mathbf{x}_i; \zeta) \quad (1)$$

$$\zeta^* := \operatorname{argmax}_{\zeta} \mathcal{L}(\zeta; \mathcal{D}) = \operatorname{argmax}_{\zeta} \prod_{i=1}^N \mathcal{P}_{Y|X}(\mathbf{y}_i; \{\hat{\mathbf{y}}_i, \hat{\nu}_i \dots \hat{\rho}_i\}) \quad (2)$$

$$\text{Uncertainty}(\hat{\mathbf{y}}_i) = \mathcal{F}(\hat{\nu}_i \dots \hat{\rho}_i) \quad (3)$$

It is common to use a *heteroscedastic* Gaussian distribution for $\mathcal{P}_{Y|X}$ [27,66], in which case $\Psi_s(\cdot; \zeta)$ is designed to predict the *mean* and *variance* of the Gaussian distribution, i.e., $\{\hat{\mathbf{y}}_i, \hat{\sigma}_i^2\} := \Psi_s(\mathbf{x}_i; \zeta)$, and the predicted *variance* itself can be treated as uncertainty in the prediction. The optimization problem becomes,

$$\zeta^* = \operatorname{argmax}_{\zeta} \prod_{i=1}^N \frac{1}{\sqrt{2\pi\hat{\sigma}_i^2}} e^{-\frac{|\hat{\mathbf{y}}_i - \mathbf{y}_i|^2}{2\hat{\sigma}_i^2}} = \operatorname{argmin}_{\zeta} \sum_{i=1}^N \frac{|\hat{\mathbf{y}}_i - \mathbf{y}_i|^2}{2\hat{\sigma}_i^2} + \frac{\log(\hat{\sigma}_i^2)}{2} \quad (4)$$

$$\text{Uncertainty}(\hat{\mathbf{y}}_i) = \hat{\sigma}_i^2. \quad (5)$$

The above equation models the per-pixel residual (between the prediction and the ground-truth) as a Gaussian distribution. However, this may not always be fit, especially in the presence of outliers and artefacts, where the residuals often follow heavy-tailed distributions. Recent works such as [61,63] have shown that heavy-tailed distributions can be modeled as a heteroscedastic generalized Gaussian distribution, in which case $\Psi_s(\cdot; \zeta)$ is designed to predict the *mean* ($\hat{\mathbf{y}}_i$), *scale* ($\hat{\alpha}_i$), and *shape* ($\hat{\beta}_i$) as trainable parameters, i.e., $\{\hat{\mathbf{y}}_i, \hat{\alpha}_i, \hat{\beta}_i\} := \Psi_s(\mathbf{x}_i; \zeta)$,

$$\zeta^* := \operatorname{argmax}_{\zeta} \mathcal{L}(\zeta) = \operatorname{argmax}_{\zeta} \prod_{i=1}^N \frac{\hat{\beta}_i}{2\hat{\alpha}_i \Gamma(\frac{1}{\hat{\beta}_i})} e^{-(|\hat{\mathbf{y}}_i - \mathbf{y}_i|/\hat{\alpha}_i)^{\hat{\beta}_i}} = \operatorname{argmin}_{\zeta} -\log \mathcal{L}(\zeta)$$

$$= \operatorname{argmin}_{\zeta} \sum_{i=1}^N \left(\frac{|\hat{\mathbf{y}}_i - \mathbf{y}_i|}{\hat{\alpha}_i} \right)^{\hat{\beta}_i} - \log \frac{\hat{\beta}_i}{\hat{\alpha}_i} + \log \Gamma\left(\frac{1}{\hat{\beta}_i}\right) \quad (6)$$

$$\text{Uncertainty}(\hat{\mathbf{y}}_i) = \frac{\hat{\alpha}_i^2 \Gamma(\frac{3}{\hat{\beta}_i})}{\Gamma(\frac{1}{\hat{\beta}_i})}. \quad (7)$$

Here $\Gamma(z) = \int_0^{\infty} x^{z-1} e^{-x} dx$, $\forall z > 0$, represents the Gamma function [2]. While the above formulation (Eq. (4)-(7)) shows the dependence of various predicted distribution parameters on one another when maximizing the likelihood, it requires training the model from scratch, that we want to avoid. In the following, we describe how we address this problem through our BayesCap.

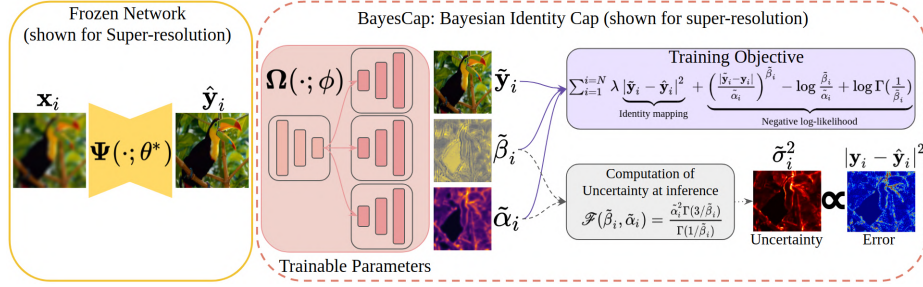


Fig. 2: BayesCap ($\Omega(\cdot; \phi)$) in tandem with the pretrained network with frozen parameters ($\Psi(\cdot; \theta^*)$) (details in Section 3.3). While the pretrained network cannot estimate the uncertainty, the proposed BayesCap feeds on the output of the pretrained network and maps it to the underlying probability distribution that allows computation of well calibrated uncertainty estimates.

3.3 Constructing BayesCap

In the above, $\Psi_s(\cdot; \zeta)$ was trained from scratch to predict all the parameters of distribution and does *not* leverage the *frozen* model $\Psi(\cdot; \theta^*)$ estimating \mathbf{y}_i using $\hat{\mathbf{y}}_i$ in a deterministic fashion. To circumvent the training from scratch, we notice that one only needs to estimate the remaining parameters of the underlying distribution. Therefore, to augment the frozen point estimation model, we learn a Bayesian identity mapping represented by $\Omega(\cdot; \phi) : \mathbb{R}^n \rightarrow \mathbb{R}^n$, that reconstructs the output of the frozen model $\Psi(\cdot; \theta^*)$ and also produces the parameters of the distribution modeling the reconstructed output. We refer to this network as BayesCap (schematic in Figure 2). As in Eq. (7), we use heteroscedastic generalized Gaussian to model output distribution, i.e.,

$$\Omega(\hat{\mathbf{y}}_i = \Psi(\mathbf{x}_i; \theta^*); \phi) = \{\tilde{\mathbf{y}}_i, \tilde{\alpha}_i, \tilde{\beta}_i\}, \text{ with } \mathbf{y}_i \sim \frac{\tilde{\beta}_i}{2\tilde{\alpha}_i \Gamma(\frac{1}{\tilde{\beta}_i})} e^{-(|\tilde{\mathbf{y}}_i - \mathbf{y}_i|/\tilde{\alpha}_i)^{\tilde{\beta}_i}} \quad (8)$$

To enforce the identity mapping, for every input \mathbf{x}_i , we regress the reconstructed output of the BayesCap ($\tilde{\mathbf{y}}_i$) with the output of the pretrained base network ($\hat{\mathbf{y}}_i$). This ensures that, the distribution predicted by BayesCap for an input \mathbf{x}_i , i.e., $\Omega(\Psi(\mathbf{x}_i; \theta^*); \phi)$, is such that the point estimates $\tilde{\mathbf{y}}_i$ match the point estimates of the pretrained network $\hat{\mathbf{y}}_i$. Therefore, as the quality of the reconstructed output improves, the uncertainty estimated by $\Omega(\cdot; \phi)$ also approximates the uncertainty for the prediction made by the pretrained $\Psi(\cdot; \theta^*)$, i.e.,

$$\tilde{\mathbf{y}}_i \rightarrow \hat{\mathbf{y}}_i \implies \tilde{\sigma}_i^2 = \frac{\tilde{\alpha}_i^2 \Gamma(3/\tilde{\beta}_i)}{\Gamma(1/\tilde{\beta}_i)} \rightarrow \hat{\sigma}_i^2 \quad (9)$$

To train $\Omega(\cdot; \phi)$ and obtain optimal parameters (ϕ^*), we minimize the fidelity term between $\tilde{\mathbf{y}}_i$ and $\hat{\mathbf{y}}_i$, along with the negative log-likelihood for $\Omega(\cdot; \phi)$, i.e.,

$$\phi^* = \underset{\phi}{\operatorname{argmin}} \sum_{i=1}^N \lambda \underbrace{|\tilde{\mathbf{y}}_i - \hat{\mathbf{y}}_i|^2}_{\text{Identity mapping}} + \underbrace{\left(\frac{|\tilde{\mathbf{y}}_i - \mathbf{y}_i|}{\tilde{\alpha}_i}\right)^{\tilde{\beta}_i} - \log \frac{\tilde{\beta}_i}{\tilde{\alpha}_i} + \log \Gamma\left(\frac{1}{\tilde{\beta}_i}\right)}_{\text{Negative log-likelihood}} \quad (10)$$

Here λ represents the hyperparameter controlling the contribution of the fidelity term in the overall loss function. Extremely high λ will lead to improper estimation of the $(\tilde{\alpha})$ and $(\tilde{\beta})$ parameters as other terms are ignored. Eq. (10) allows BayesCap to estimate the underlying distribution and uncertainty.

Eq. (8) and (10) show that the construction of Ω is independent of Ψ and the Ω always performs the Bayesian identity mapping regardless of the task performed by Ψ . This suggests that task specific tuning will have minimal impact on the performance of Ω . In our experiments, we employed the same network architecture for BayesCap with the same set of hyperparameters, across different tasks and it achieves state-of-the-art uncertainty estimation results without task specific tuning (as shown in Section 4), highlighting that BayesCap is not sensitive towards various design choices including architecture, learning-rate, etc.

4 Experiments

We first describe our experimental setup (i.e., datasets, tasks, and evaluation metrics) in Section 4.1. We compare our model to a wide variety of state-of-the-art methods quantitatively and qualitatively in Section 4.2. Finally in Section 4.3, we provide an ablation analysis along with a real world application of BayesCap for detecting out-of-distribution samples.

4.1 Tasks and Datasets

We show the efficacy of our BayesCap on various image enhancement and translation tasks including super-resolution, deblurring, inpainting, and MRI translation, as detailed below. In general, image enhancement and translation tasks are highly ill-posed problems as an *injective function* between input and output may not exist [39,61], thereby necessitating the need to learn a probabilistic mapping to quantify the uncertainty and indicate poor reconstruction of the output images. For each task we choose a well established deterministic pretrained network, for which we estimate uncertainties.

Super-resolution. The goal is to map low-resolution images to their high-resolution counterpart. We choose pretrained SRGAN [36] as our base model $\Psi(\cdot; \theta^*)$. The BayesCap model $\Omega(\cdot; \phi)$ is trained on ImageNet patches sized 84×84 to perform $4 \times$ super-resolution. The resulting combination of SRGAN and BayesCap is evaluated on the Set5 [5], Set14 [76], and BSD100 [41] datasets.

Deblurring. The goal is to remove noise from images corrupted with blind motion. We use the pretrained DeblurGANv2 [31] which shows improvements over the original DeblurGAN [30]. The BayesCap model is evaluated on the GoPro dataset [44], using standard train/test splits.

Inpainting. The goal is to fill masked regions of an input image. We use pretrained DeepFillv2 [75], that improves over DeepFill [74], as the base model for inpainting. Both the original base model and the **BayesCap** are trained and tested on the standard train/test split of Places365 dataset [79].

MRI Translation. We predict the T2 MRI from T1 MRI, an important problem in medical imaging as discussed in [8,7,24,72,73]. We use the pretrained deterministic UNet as base model [63,52]. Both the base model and **BayesCap** are trained and tested on IXI dataset [51] following [63].

Baselines. For all tasks, we compare **BayesCap** against 7 methods in total, out of which 6 baselines can estimate uncertainty of a pretrained model without re-training and one baseline modifies the base network and train it from *scratch* to estimate the uncertainty. The first set of baselines belong to *test-time data augmentation* (TTDA) technique [3,65,66], where we generate multiple perturbed copies of the input and use the set of corresponding outputs to compute the uncertainty. We consider three different ways of perturbing the input, (i) per-pixel noise perturbations **TTDAp** [3,65], (ii) affine transformations **TTDAa** [3,65] and (iii) random corruptions from Gaussian blurring, contrast enhancement, and color jittering (**TTDAc**) [3,65]. As additional baseline, we also consider **TTDApac** that generates the multiple copies by combining pixel-level perturbations, affine transformations, and corruptions as described above.

Another set of baselines uses dropout [56,26,17,34,40] before the final predictions. This is possible even for the models *that are not originally trained with dropout*. We refer to this model as **DO**. In addition, we consider a baseline that combines dropout with test-time data augmentation (**DOpac**). Finally, we also compare against a model trained from scratch to produce the uncertainty as described in [27]. We refer to this as **Scratch**.

Metrics. We evaluate the performance of various models on two kinds of metrics (i) image reconstruction quality and (ii) predicted uncertainty calibration quality. To measure reconstruction quality, we use SSIM [69] and PSNR. For inpainting we also show mean ℓ_1 and ℓ_2 error, following the convention in original works [74,75]. We emphasize that all the methods, *except Scratch*, can use the output of the pretrained base model and only derive the uncertainty maps using different estimation techniques described above. Therefore image reconstruction quality metrics like SSIM, PSNR, mean ℓ_1 and ℓ_2 error remain the same as that of base network. However, **Scratch** method *does not* have access to the pretrained model, therefore it has to use its own predicted output and uncertainty estimates.

To quantify the quality of the uncertainty, we use the *uncertainty calibration error* (UCE) as described in [35,23] for regression tasks. It measures the discrepancy between the predictive error and predictive uncertainty, given by, $\text{UCE} := \sum_{m=1}^M |B_m|/N |\text{err}(B_m) - \text{uncer}(B_m)|$, where B_m is one of the uniformly separated bins, $\text{err}(B_m) := 1/|B_m| \sum_{i \in B_m} \|\hat{\mathbf{y}}_i - \mathbf{y}_i\|^2$, and $\text{uncer}(B_m) := 1/|B_m| \sum_{i \in B_m} \hat{\sigma}_i^2$. We also use *correlation coefficient* (C.Coeff.) between the error and the uncertainty, as high correlation is desirable.

Implementation Details. We optimize Eq. (10) using the Adam optimizer [28] and a batch size of 2 with images that are resized to 256×256 .




Fig. 3: Input (LR, x) and output of pretrained SRGAN (SR, \hat{y}) along with output of BayesCap ($\{\tilde{y}, \tilde{\alpha}, \tilde{\beta}\}$). Spatially varying parameters ($\tilde{\alpha}, \tilde{\beta}$) lead to well-calibrated uncertainty $\tilde{\sigma}^2$, highly correlated with the SRGAN error, $|\hat{y} - y|^2$.




Fig. 4: Qualitative example showing the results of the pre-trained SRGAN model along with the uncertainty maps produced by BayesCap and the other methods. Uncertainty derived from BayesCap has better correlation with the error.

During training we exponentially anneal the hyperparameter λ that is initially set to 10. This guides the BayesCap to learn the identity mapping in the beginning, and gradually learn the optimal parameters of the underlying distribution via maximum likelihood estimation.

4.2 Results

Super-resolution. Table 1 shows the image reconstruction performance along with the uncertainty calibration performance on the Set5, Set14, and the BSD100 datasets for all the methods. We see that BayesCap significantly outperforms all other methods in terms of UCE and C.Coeff while retaining the image reconstruction performance of the base model. For instance, across all the 3 sets, the correlation coefficient is always greater than 0.4 showing that the error and the uncertainty are correlated to a high degree. The model trained from scratch has a correlation coefficient between 0.22-0.31 which is lower than BayesCap. The baselines based on dropout and test-time data augmentation show nearly no

D	Metrics	SRGAN	TTDA _p	TTDA _a	TTDA _c	TTDA _{pac}	DO	DO _{pac}	Scratch	BayesCap
Set5	PSNR \uparrow	29.40	29.40	29.40	29.40	29.40	29.40	29.40	27.83	29.40
	SSIM \uparrow	0.8472	0.8472	0.8472	0.8472	0.8472	0.8472	0.8472	0.8166	0.8472
	UCE \downarrow	NA	0.39	0.40	0.42	0.47	0.33	0.36	0.035	0.014
	C.Coeff \uparrow	NA	0.17	0.13	0.08	0.03	0.05	0.07	0.28	0.47
Set14	PSNR \uparrow	26.02	26.02	26.02	26.02	26.02	26.02	26.02	25.31	26.02
	SSIM \uparrow	0.7397	0.7397	0.7397	0.7397	0.7397	0.7397	0.7397	0.7162	0.7397
	UCE \downarrow	NA	0.57	0.63	0.61	0.69	0.48	0.52	0.048	0.017
	C.Coeff \uparrow	NA	0.07	0.04	0.04	0.06	0.08	0.04	0.22	0.42
BSD100	PSNR \uparrow	25.16	25.16	25.16	25.16	25.16	25.16	25.16	24.39	25.16
	SSIM \uparrow	0.6688	0.6688	0.6688	0.6688	0.6688	0.6688	0.6688	0.6297	0.6688
	UCE \downarrow	NA	0.72	0.77	0.81	0.83	0.61	0.64	0.057	0.028
	C.Coeff \uparrow	NA	0.13	0.09	0.11	0.09	0.10	0.08	0.31	0.45

Table 1: Quantitative results showing the performance of pretrained SRGAN in terms of PSNR and SSIM, along with the quality of uncertainty maps obtained by BayesCap and other baselines, in terms of UCE and Correlation Coefficient (C.Coeff). All results on 3 datasets including Set5, Set14, and BSD100.

correlation between the uncertainty estimates and the error (C.Coeff. of 0.03-0.17). A similar trend can be seen with the UCE where BayesCap has the best UCE scores followed by the model trained from scratch, while the test-time data augmentation and the dropout baselines have a very high UCE (between 0.33 and 0.83) suggesting poorly calibrated uncertainty estimates.

Qualitatively, Figure 3 shows the prediction of the pretrained SRGAN along with the predictions of the BayesCap showing per-pixel estimated distribution parameters along with the uncertainty map on a sample from Set5 dataset. High correlation between the per-pixel predictive error of SRGAN and uncertainties from BayesCap suggests that BayesCap produces well-calibrated uncertainty estimates. Moreover, Figure 4 shows that uncertainty produced by other baselines are not in agreement with the error (e.g., TTDA_p does not show high uncertainty within the eye, where error is high) indicating that they are poorly calibrated.

Deblurring. We report the results on the GoPro dataset in Table 2. DeblurGANv2 achieves significantly better results than Scratch (29.55 vs 26.16 PSNR). In terms of UCE, BayesCap outperforms all the methods by achieving a low score of 0.038. While the Scratch is close and achieves a UCE of 0.076, all the other methods have a UCE that is nearly 10 times higher suggesting that BayesCap estimates the most calibrated uncertainty. This is also visible in Figure 5-(left) where the uncertainties provided by BayesCap is correlated with the error (C.Coeff. of 0.32) unlike methods from TTDA and DO class that have very low correlation between the uncertainty and the error (C.Coeff of 0.03 - 0.17). While Scratch achieves a reasonable score, second only to BayesCap, in terms of UCE (0.076 vs. 0.038) and C.Coeff (0.21 vs. 0.32), it has much poorer image reconstruction output with a PSNR of 26.16 and SSIM of 0.8136. The poor reconstruction also justifies the relatively higher uncertainty values for Scratch (as seen in Figure 5-(left)) when compared to BayesCap.

Inpainting. Table 3 shows the results on Places365 dataset. The pretrained base model DeepFillv2 [75] achieves a mean L1 error of 9.1%, however Scratch




Fig. 5: Qualitative example showing the results of the pretrained DeblurGANv2 and DeepFillv2 on image deblurring (left) and inpainting (right) tasks along with the uncertainty maps produced by different methods.

D	Metrics	DeblurGANv2	TTDAp	TTDAa	TTDAc	TTDApac	DO	DOpac	Scratch	BayesCap
GoPro	PSNR \uparrow	29.55	29.55	29.55	29.55	29.55	29.55	29.55	26.16	29.55
	SSIM \uparrow	0.9340	0.9340	0.9340	0.9340	0.9340	0.9340	0.9340	0.8136	0.9340
	UCE \downarrow	NA	0.44	0.45	0.49	0.53	0.52	0.59	0.076	0.038
	C.Coeff \uparrow	NA	0.17	0.13	0.08	0.03	0.05	0.07	0.21	0.32

Table 2: Results showing the performance of pretrained DeblurGANv2 in terms of PSNR and SSIM, along with the quality of uncertainty obtained by BayesCap and other methods, in terms of UCE and C.Coeff on GoPro dataset.

D	Metrics	DeepFillv2	TTDAp	TTDAa	TTDAc	TTDApac	DO	DOpac	Scratch	BayesCap
Places365	m. ℓ_1 err. \downarrow	9.1%	9.1%	9.1%	9.1%	9.1%	9.1%	9.1%	15.7%	9.1%
	m. ℓ_2 err. \downarrow	1.6%	1.6%	1.6%	1.6%	1.6%	1.6%	1.6%	5.8%	1.6%
	PSNR \uparrow	18.34	18.34	18.34	18.34	18.34	18.34	18.34	17.24	18.34
	SSIM \uparrow	0.6285	0.6285	0.6285	0.6285	0.6285	0.6285	0.6285	0.6032	0.6285
	UCE \downarrow	NA	0.63	0.88	0.87	0.93	1.62	1.49	0.059	0.011
	C.Coeff \uparrow	NA	0.26	0.11	0.12	0.08	0.09	0.12	0.44	0.68

Table 3: Performance of pretrained DeepFillv2 in terms of mean ℓ_1 error, mean ℓ_2 error, PSNR and SSIM, along with the quality of uncertainty obtained by BayesCap and other methods, in terms of UCE and C.Coeff on Places365 dataset.

is much worse, achieving a mean L1 error of 15.7%. This again demonstrates that training a Bayesian model from scratch often does not replicate the performance of deterministic counterparts. Also, BayesCap retains the reconstruction performance of DeepFillv2 [75] and provides well-calibrated uncertainties, as demonstrated by highest C.Coeff. (0.68), and the lowest UCE (0.011). Methods belonging to TTDA and DO classes are unable to provide good uncertainties (C.Coeff. of 0.08-0.26). The example in Figure 5-(right) also illustrates an interesting phenomenon. Although, the uncertainties are predicted for the entire image, we see that BayesCap automatically learns to have extremely low uncertainty values outside the masked region which is perfectly reconstructed. Within the masked region, uncertainty estimates are highly correlated with the error.

MRI Translation. We perform T1 to T2 MRI translation as described in [63,61] which has an impact in clinical settings by reducing MRI acquisition times [8,7,24,72,73]. The quantitative results on the IXI dataset are show in Table 4. The pretrained base model employing U-Net architecture, achieves a SSIM score of 0.9272 which is nearly matched by Scratch (0.9169). However, we see that BayesCap performs better than Scratch in terms of UCE (0.036

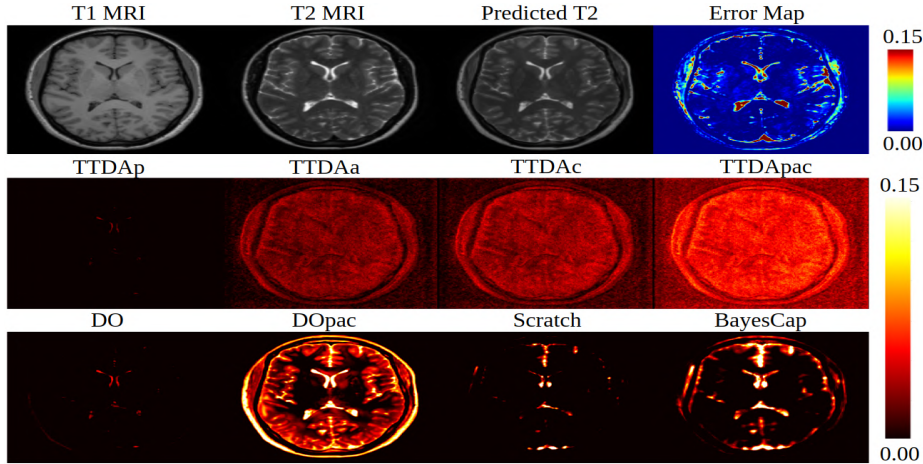


Fig. 6: Qualitative example showing the results of the pretrained UNet for T1 to T2 MRI translation along with the uncertainty produced by different methods.

D	Metrics	UNet	TTDAp	TTDAa	TTDAc	TTDApac	DO	DOpac	Scratch	BayesCap
IXI	PSNR \uparrow	25.70	25.70	25.70	25.70	25.70	25.70	25.70	25.50	25.70
	SSIM \uparrow	0.9272	0.9272	0.9272	0.9272	0.9272	0.9272	0.9272	0.9169	0.9272
	UCE \downarrow	NA	0.53	0.46	0.41	0.44	0.38	0.40	0.036	0.029
	C.Coeff \uparrow	NA	0.05	0.14	0.16	0.08	0.13	0.47	0.52	0.58

Table 4: Performance of pretrained UNet for MRI translation in terms of PSNR and SSIM, along with the quality of uncertainty obtained by **BayesCap** and other methods, in terms of UCE and C.Coeff on IXI Dataset.

vs. 0.029) and C.Coeff (0.52 vs. 0.58). Other methods are poorly calibrated, as indicated by high UCE and low C.Coeff. This is also evident from Figure 6, indicating high correlation between the **BayesCap** uncertainty and error, while low correlation of the same for other methods.

4.3 Ablation Studies

As discussed in Section 3.3, **BayesCap** can help in estimating uncertainty for the frozen model only if **BayesCap** provides perfect reconstruction. To study this, we design an experiment that deteriorates the identity mapping learned by **BayesCap**, leading to poor uncertainty estimates. Moreover, we also demonstrate that **BayesCap** is much more data and compute efficient when compared to **Scratch** (i.e., model capable of estimating the uncertainty, trained from scratch).

Figure 7 shows that preserving the identity mapping is essential to providing well-calibrated post-hoc uncertainty estimates using **BayesCap**. Here, we gradually feed increasingly noisy samples (corrupted by zero-mean Gaussian with variance given by κ^2) to the model that leads to poor reconstruction by **BayesCap** and as a result degrades the identity mapping. As the quality of the identity mapping degrades (decreasing SSIM), we see that quality of uncertainty also degrades (increasing UCE). For instance, for super-resolution, with zero noise in the input the reconstruction quality of the **BayesCap** is at its peak (SSIM, 0.8472)




Fig. 7: *Impact of the identity mapping.* Degrading the quality of the identity mapping (SSIM) at inference, leads to poorly calibrated uncertainty (UCE). κ represents the magnitude of noise used for degrading the identity mapping.




Fig. 8: BayesCap can be trained to achieve optimal performance in fewer epochs (left), while being more data-efficient (achieves better results with fewer samples) as compared to Scratch (middle and right), shown for super-resolution.

leading to almost identical mapping. This results in well-calibrated uncertainty maps derived from BayesCap (UCE, 0.014). However, with $\kappa = 0.15$, the reconstruction quality decreases sharply (SSIM of 0.204) leading to poorly calibrated uncertainty (UCE of 0.0587), justifying the need for the identity mapping.

We also show that BayesCap is more efficient than training a Bayesian model from scratch both in terms of time and data required to train the model in Figure 8-(Left). On all the 4 tasks, BayesCap can be trained 3-5 \times faster than Scratch. For super-resolution, we show that BayesCap can achieve competitive results even when trained on a fraction of the dataset, whereas the performance of Scratch reduces sharply in low data regime, as shown in Figure 8-(Middle and Right). For instance, with just 50% of training data, BayesCap performs 33% better than Scratch in terms of SSIM. This is because BayesCap learns the autoencoding task, whereas Scratch learns the original translation task.

4.4 Application: Out-of-Distribution Analysis

The proposed BayesCap focuses on estimating well-calibrated uncertainties of pretrained image regression models in a post-hoc fashion. This can be crucial in many critical real-world scenarios such as autonomous driving and navigation [71,27,68,43]. We consider monocular depth estimation (essential for autonomous driving) and show that the estimated uncertainty can help in detecting *out-of-distribution* (OOD) samples. We take MonoDepth2 [21] that is trained on the KITTI dataset [19]. The resulting model with BayesCap is evaluated on the validation sets of the KITTI dataset, the India Driving Dataset (IDD) [64] as well as the Places365 dataset [79]. The KITTI dataset captures images from German cities, while the India Driving Dataset has images from Indian roads. Places365 consists of images for scene recognition and is vastly different from driving datasets. Intuitively, both IDD and Places365 represent OOD as there is a shift in distribution when compared to training data (KITTI), this is captured




Fig. 9: BayesCap with MonoDepth2 [21] for depth estimation in autonomous driving. Trained on KITTI and evaluated on (a) KITTI, (b) Indian Driving Dataset, and (c) Places365. (d) and (e) Plots show mean uncertainty values and ROC curve for OOD detection respectively, as described in Section 4.4.

in Figure 9-(a,b,c,d), representing degraded depth and uncertainty on OOD images. The uncertainties also reflect this with increasingly higher values for OOD samples as also shown in the bar plot (Figure 9-(d)). This suggests that mean uncertainty value for an image can be used to detect if it belongs to OOD.

To quantify this, we plot the ROC curve for OOD detection on a combination of the KITTI (in-distribution) samples and IDD and Places365 (out-of-distribution) samples in Figure 9-(e). Additionally, we compare using uncertainty for OOD detection against (i) using intermediate features from the pretrained model [49] and (ii) using features from the bottleneck of a trained autoencoder [80]. Samples are marked OOD if the distance between the features of the samples and mean feature of the KITTI validation set is greater than a threshold. We clearly see that using the mean uncertainty achieves far superior results (AUROC of 0.96), against the pretrained features (AUROC of 0.82) and autoencoder based approach (AUROC of 0.72). Despite not being specifically tailored for OOD detection, BayesCap achieves strong results. This indicates the benefits of its well-calibrated uncertainty estimates on downstream tasks.

5 Conclusions

We proposed BayesCap, a fast and cheap post-hoc uncertainty estimation method for pretrained deterministic models. We show that our method consistently produces well-calibrated uncertainty estimates across a wide variety of image enhancement and translation tasks without hampering the performance of the pretrained model. This is in sharp contrast to training a Bayesian model from scratch that is more expensive and often not competitive with deterministic counterparts. We demonstrate that derived calibrated uncertainties can be used in critical scenarios for detecting OOD and helping decision-making. One limitation of our model is the assumption that the input is sufficiently contained in the output of the target task. Future works may address this limitation, as well as extending BayesCap to discrete predictions.

Acknowledgments. This work has been partially funded by the ERC (853489 - DEXIM), by the DFG (2064/1 - Project number 390727645). The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Uddeshya Upadhyay and Shyamgopal Karthik.

References

1. Armanious, K., Jiang, C., Fischer, M., Küstner, T., Hepp, T., Nikolaou, K., Gatidis, S., Yang, B.: Medgan: Medical image translation using gans. *Computerized medical imaging and graphics* (2020)
2. Artin, E.: *The gamma function*. Courier Dover Publications (2015)
3. Ayhan, M.S., Berens, P.: Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. *MIDL* (2018)
4. Bae, G., Budvytis, I., Cipolla, R.: Estimating and exploiting the aleatoric uncertainty in surface normal estimation. In: *IEEE ICCV* (2021)
5. Bevilacqua, M., Roumy, A., Guillemot, C., Alberi-Morel, M.L.: Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In: *BMVC* (2012)
6. Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural network. In: *ICML* (2015)
7. Bowles, C., Qin, C., Ledig, C., Guerrero, R., Gunn, R., Hammers, A., Sakka, E., Dickie, D.A., Hernández, M.V., Royle, N., et al.: Pseudo-healthy image synthesis for white matter lesion segmentation. In: *International Workshop on Simulation and Synthesis in Medical Imaging* (2016)
8. Chartsias, A., Joyce, T., Giuffrida, M.V., Tsafaris, S.A.: Multimodal mr synthesis via modality-invariant latent representation. *IEEE TMI* (2017)
9. Chen, T., Fox, E., Guestrin, C.: Stochastic gradient hamiltonian monte carlo. In: *ICML*. PMLR (2014)
10. Coglianese, C., Lehr, D.: Regulating by robot: Administrative decision making in the machine-learning era. *Geo. LJ* (2016)
11. Cohen, J.P., Luck, M., Honari, S.: Distribution matching losses can hallucinate features in medical image translation. In: *MICCAI* (2018)
12. Daxberger, E., Kristiadi, A., Immer, A., Eschenhagen, R., Bauer, M., Hennig, P.: Laplace redux-effortless bayesian deep learning. *NeurIPS* (2021)
13. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. *IEEE TPAMI* (2015)
14. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: Flownet: Learning optical flow with convolutional networks. In: *IEEE ICCV* (2015)
15. Eschenhagen, R., Daxberger, E., Hennig, P., Kristiadi, A.: Mixtures of Laplace approximations for improved *Post-Hoc* uncertainty in deep learning. *NeurIPS Workshop on Bayesian Deep Learning* (2021)
16. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: *IEEE CVPR* (2018)
17. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *ICML* (2016)
18. Gawlikowski, J., Tassi, C.R.N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., et al.: A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342* (2021)
19. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* (2013)
20. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: *IEEE CVPR* (2017)
21. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: *IEEE ICCV* (2019)

22. Graves, A.: Practical variational inference for neural networks. NIPS (2011)
23. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: ICML. PMLR (2017)
24. Iglesias, J.E., Konukoglu, E., Zikic, D., Glocker, B., Leemput, K.V., Fischl, B.: Is synthesizing mri contrast useful for inter-modality analysis? In: MICCAI (2013)
25. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Let there be color! joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. ACM Transactions on Graphics (ToG) (2016)
26. Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. arXiv preprint arXiv:1511.02680 (2015)
27. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? NIPS (2017)
28. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. ICLR (2015)
29. Kuleshov, V., Fenner, N., Ermon, S.: Accurate uncertainties for deep learning using calibrated regression. In: ICML (2018)
30. Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J.: Deblurgan: Blind motion deblurring using conditional adversarial networks. In: IEEE CVPR (2018)
31. Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In: IEEE ICCV (2019)
32. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. arXiv preprint arXiv:1612.01474 (2016)
33. Laves, M.H., Ihler, S., Fast, J.F., Kahrs, L.A., Ortmaier, T.: Well-calibrated regression uncertainty in medical imaging with deep learning. In: MIDL (2020)
34. Laves, M.H., Ihler, S., Kortmann, K.P., Ortmaier, T.: Well-calibrated model uncertainty with temperature scaling for dropout variational inference. arXiv preprint arXiv:1909.13550 (2019)
35. Laves, M.H., Ihler, S., Kortmann, K.P., Ortmaier, T.: Calibration of model uncertainty for dropout variational inference. arXiv preprint arXiv:2006.11584 (2020)
36. Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., Shi, W.: Photo-realistic single image super-resolution using a generative adversarial network. In: IEEE CVPR (2017)
37. Levi, D., Gispan, L., Giladi, N., Fetaya, E.: Evaluating and calibrating uncertainty prediction in regression tasks. arXiv preprint arXiv:1905.11659 (2019)
38. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. Medical image analysis (2017)
39. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. NIPS (2017)
40. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for bayesian uncertainty in deep learning. NeurIPS (2019)
41. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: IEEE ICCV (2001)
42. McAllister, R., Gal, Y., Kendall, A., Van Der Wilk, M., Shah, A., Cipolla, R., Weller, A.: Concrete problems for autonomous vehicle safety: Advantages of bayesian deep learning. In: IJCAI (2017)
43. Michelmore, R., Kwiatkowska, M., Gal, Y.: Evaluating uncertainty quantification in end-to-end autonomous driving control. arXiv preprint arXiv:1811.06817 (2018)
44. Nah, S., Hyun Kim, T., Mu Lee, K.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: IEEE CVPR (2017)

45. Osawa, K., Swaroop, S., Khan, M.E.E., Jain, A., Eschenhagen, R., Turner, R.E., Yokota, R.: Practical deep learning with bayesian principles. *NeurIPS* (2019)
46. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: *IEEE CVPR* (2016)
47. Phan, B., Salay, R., Czarnecki, K., Abdelzad, V., Denouden, T., Vernekar, S.: Calibrating uncertainties in object localization task. *arXiv preprint arXiv:1811.11210* (2018)
48. Plotz, T., Roth, S.: Benchmarking denoising algorithms with real photographs. In: *IEEE CVPR* (2017)
49. Reiss, T., Cohen, N., Bergman, L., Hoshen, Y.: Panda: Adapting pretrained features for anomaly detection and segmentation. In: *IEEE CVPR* (2021)
50. Riquelme, C., Tucker, G., Snoek, J.: Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *arXiv preprint arXiv:1802.09127* (2018)
51. Robinson, E.C., Hammers, A., Ericsson, A., Edwards, A.D., Rueckert, D.: Identifying population differences in whole-brain structural networks: a machine learning approach. *NeuroImage* (2010)
52. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *MICCAI* (2015)
53. van der Schaar, M., Alaa, A.M., Floto, A., Gimson, A., Scholtes, S., Wood, A., McKinney, E., Jarrett, D., Lio, P., Ercole, A.: How artificial intelligence and machine learning can help healthcare systems respond to covid-19. *Machine Learning* (2021)
54. Schwarting, W., Alonso-Mora, J., Rus, D.: Planning and decision-making for autonomous vehicles. *Annual Review of Control, Robotics, and Autonomous Systems* (2018)
55. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *ICLR* (2015)
56. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *JMLR* (2014)
57. Sudarshan, V.P., Upadhyay, U., Egan, G.F., Chen, Z., Awate, S.P.: Towards lower-dose pet using physics-based uncertainty-aware multimodal learning with robustness to out-of-distribution data. *Medical Image Analysis* (2021)
58. Tian, C., Fei, L., Zheng, W., Xu, Y., Zuo, W., Lin, C.W.: Deep learning on image denoising: An overview. *Neural Networks* (2020)
59. Upadhyay, U., Awate, S.P.: A mixed-supervision multilevel gan framework for image quality enhancement. In: *MICCAI* (2019)
60. Upadhyay, U., Awate, S.P.: Robust super-resolution gan, with manifold-based and perception loss. In: *IEEE International Symposium on Biomedical Imaging* (2019)
61. Upadhyay, U., Chen, Y., Akata, Z.: Robustness via uncertainty-aware cycle consistency. *NeurIPS* (2021)
62. Upadhyay, U., Chen, Y., Hepp, T., Gatidis, S., Akata, Z.: Uncertainty-guided progressive gans for medical image translation. In: *MICCAI* (2021)
63. Upadhyay, U., Sudarshan, V.P., Awate, S.P.: Uncertainty-aware gan with adaptive loss for robust mri image enhancement. In: *IEEE ICCV Workshop* (2021)
64. Varma, G., Subramanian, A., Namboodiri, A., Chandraker, M., Jawahar, C.: Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In: *IEEE WACV* (2019)
65. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Test-time augmentation with uncertainty estimation for deep learning-based medical image segmentation. *MIDL* (2018)

66. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing* (2019)
67. Wang, X., Aitchison, L.: Bayesian ood detection with aleatoric uncertainty and outlier exposure. In: *Fourth Symposium on Advances in Approximate Bayesian Inference* (2021)
68. Wang, Y., Chao, W.L., Garg, D., Hariharan, B., Campbell, M., Weinberger, K.Q.: Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In: *IEEE CVPR* (2019)
69. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE TIP* (2004)
70. Welling, M., Teh, Y.W.: Bayesian learning via stochastic gradient langevin dynamics. In: *ICML* (2011)
71. Xu, W., Pan, J., Wei, J., Dolan, J.M.: Motion planning under uncertainty for on-road autonomous driving. In: *IEEE ICRA* (2014)
72. Ye, D.H., Zikic, D., Glocker, B., Criminisi, A., Konukoglu, E.: Modality propagation: coherent synthesis of subject-specific scans with data-driven regularization. In: *MICCAI* (2013)
73. Yu, B., Zhou, L., Wang, L., Shi, Y., Fripp, J., Bourgeat, P.: Ea-gans: edge-aware generative adversarial networks for cross-modality mr image synthesis. *IEEE TMI* (2019)
74. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: *IEEE CVPR* (2018)
75. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. In: *IEEE CVPR* (2019)
76. Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparse-representations. In: *International conference on curves and surfaces* (2010)
77. Zhang, J., Kailkhura, B., Han, T.Y.J.: Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In: *ICML* (2020)
78. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: *ECCV* (2016)
79. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE TPAMI* (2017)
80. Zhou, C., Paffenroth, R.C.: Anomaly detection with robust deep autoencoders. In: *ACM KDD* (2017)
81. Zhu, Y., Tang, Y., Tang, Y., Elton, D.C., Lee, S., Pickhardt, P.J., Summers, R.M.: Cross-domain medical image translation by shared latent gaussian mixture model. In: *MICCAI* (2020)