

Research Article

# A Critical Synthesis of Uncertainty Quantification and Foundation Models in Monocular Depth Estimation

Steven Landgraf<sup>1</sup>, R. Qin<sup>2</sup>, Markus Ulrich<sup>1</sup>

1. Institute of Photogrammetry and Remote Sensing (IPF), Karlsruhe Institute of Technology, Karlsruhe, Germany; 2. The Ohio State University, USA

While recent foundation models have enabled significant breakthroughs in monocular depth estimation, a clear path towards safe and reliable deployment in the real-world remains elusive. Metric depth estimation, which involves predicting absolute distances, poses particular challenges, as even the most advanced foundation models remain prone to critical errors. Since quantifying the uncertainty has emerged as a promising endeavor to address these limitations and enable trustworthy deployment, we fuse five different uncertainty quantification methods with the current state-of-the-art DepthAnythingV2 foundation model. To cover a wide range of metric depth domains, we evaluate their performance on four diverse datasets. Our findings identify fine-tuning with the Gaussian Negative Log-Likelihood Loss (GNLL) as a particularly promising approach, offering reliable uncertainty estimates while maintaining predictive performance and computational efficiency on par with the baseline, encompassing both training and inference time. By fusing uncertainty quantification and foundation models within the context of monocular depth estimation, this paper lays a critical foundation for future research aimed at improving not only model performance but also its explainability. Extending this critical synthesis of uncertainty quantification and foundation models into other crucial tasks, such as semantic segmentation and pose estimation, presents exciting opportunities for safer and more reliable machine vision systems.

Corresponding author: Steven Landgraf, [steven.landgraf@kit.edu](mailto:steven.landgraf@kit.edu)

# 1. Introduction

Monocular depth estimation (MDE) received significant attention in recent years due to its crucial role in various downstream tasks ranging from autonomous driving<sup>[1][2]</sup> and robotics<sup>[3][4]</sup> to AI-generated content such as images<sup>[5]</sup>, videos<sup>[6]</sup>, and 3D scenes<sup>[7][8][9]</sup>. At its core, MDE aims to transform a single image into a depth map by regressing range values for each pixel, all without exploiting direct range or stereo measurements. Theoretically, MDE is a geometrically ill-posed problem that is fundamentally ambiguous and can only be solved with the help of prior knowledge about object shapes, sizes, scene layouts, and occlusion patterns. This inherent requirement for scene understanding perfectly aligns MDE with deep learning approaches, which have proven proficient in encoding potent priors<sup>[10][11][12]</sup><sup>[13]</sup>. These models benefit from extreme scaling, i.e., training on massive datasets and increasing model size, which facilitates the emergence of high-level visual scene understanding.

Based on these findings, a plethora of models have been proposed to address the challenges of MDE, with recent state-of-the-art solutions often leveraging large vision transformers trained on internet-scale data<sup>[14][15][16][17][10][11][12][13]</sup>, yielding foundation models capable of generalizing to a wide range of applications and scenes. A particularly challenging yet crucial application in fields such as robotics<sup>[3]</sup><sup>[4]</sup>, augmented reality<sup>[18]</sup>, and autonomous driving<sup>[1][2]</sup> is the estimation of absolute distances in real-world units (e.g., meters), commonly referred to as metric depth estimation. This task is especially difficult due to inherent metric ambiguities caused by different camera models and scene variations. Fortunately, these foundation models can successfully be fine-tuned in the respective domain<sup>[10][11][12]</sup><sup>[13]</sup> to determine exceptionally accurate metric depths.

However, the strong performance of foundation models on common benchmarks<sup>[19][20][21][22]</sup> can lead to naive deployment, potentially overlooking their limitations. This can be particularly detrimental in safety-critical applications where errors can have serious consequences. There are multiple challenges associated with real-world deployment of deep learning models, including the lack of transparency due to the "black box" character of end-to-end systems<sup>[23][24]</sup>, the inability to distinguish between in-domain and out-of-domain samples<sup>[25][26]</sup>, the tendency to be overconfident<sup>[27]</sup>, and the sensitivity to adversarial attacks<sup>[28][29][30]</sup>.

To mitigate these risks, recent research has advocated for the quantification of uncertainty in deep learning models<sup>[31][32][33][34][35][36][37]</sup>, particularly in scenarios where their deployment could have real-world implications. A number of promising methods for uncertainty quantification (UQ) have

been developed, yet, surprisingly, there has been little attention on the integration of UQ with powerful MDE foundation models. As illustrated in Figure 1, even state-of-the-art foundation models are not immune to inaccuracies. However, if we can leverage UQ to correlate high uncertainties with erroneous predictions, it opens up the possibility of safer deployment of these models in real-world applications. To bridge the gap between ground-breaking results in research and safe, reliable deployment in real-world applications, we investigate multiple UQ methods in combination with MDE foundation models. We specifically focus on combining the state-of-the-art DepthAnythingV2 foundation model<sup>[43]</sup> with five different UQ methods to enable pixel-wise uncertainty measures for metric depth estimation:

1. **Learned Confidence (LC)**<sup>[38]</sup>: Confidences, interpreted as uncertainties, are learned by extending the primary objective function with an additional loss term.
2. **Gaussian Negative Log-Likelihood (GNLL)**<sup>[39]</sup>: Predictions are treated as samples from a Gaussian distribution, with the network outputting both a predictive mean and its corresponding variance, which is learned implicitly through minimizing the Gaussian Negative Log-Likelihood.
3. **Monte Carlo Dropout (MCD)**<sup>[40]</sup>: Dropout layers remain active during inference, sampling from the posterior distribution to estimate a predictive mean and variance.
4. **Sub-Ensembles (SE)**<sup>[41]</sup>: A Deep Ensemble<sup>[42]</sup> is approximated by multiplying a subset of the model's layers instead of using the full model.
5. **Test-Time Augmentation (TTA)**<sup>[43]</sup>: Perturbations applied to inputs during inference produce unique samples, enabling computation of a predictive mean and corresponding variance.


We conduct extensive evaluations across four diverse datasets:

1. **NYUv2**<sup>[19]</sup>: Low-resolution indoor scenes.
2. **Cityscapes**<sup>[21]</sup>: High-resolution outdoor urban scenes.
3. **HOPE**<sup>[44]</sup>: Synthetic training images and real-world testing images of household objects.
4. **UseGeo**<sup>[45]</sup>: High-resolution aerial images.

This wide selection ensures coverage of a broad spectrum of real-world applications. Additionally, we examine all publicly available pre-trained encoders of the DepthAnythingV2 model, comprehensively revealing insights into the performance of models across various sizes and configurations.

We believe that our analysis will contribute to a deeper understanding of the limitations of these powerful foundation models, enhancing transparency and reliability in MDE. By highlighting often

overlooked challenges, such as the lack of explainability, our work paves the way for future research that not only seeks to improve performance but also aims to teach these models to know what they don't know.



**Figure 1.** Qualitative example of a fine-tuned DepthAnythingV2 for metric monocular depth estimation on the NYUv2 dataset [19], using a ViT-S encoder and Monte Carlo Dropout for an additional uncertainty estimate. The binary accuracy map is based on the  $\delta_1$  error. The strong correlation between erroneous predictions and high uncertainties highlights the potential of integrating uncertainty quantification (UQ) methods with foundation models for MDE.

## 2. Related Work

### 2.1. Monocular Depth Estimation

**Foundations.** Monocular Depth Estimation (MDE) is a dense regression task that aims to predict a depth value for each pixel in a given input image. The pioneering work of<sup>[46]</sup> laid the foundation for MDE by directly predicting depth using a multi-scale neural network. This seminal approach demonstrated that convolutional neural networks could effectively learn spatial hierarchies and capture depth cues from monocular images, thus inspiring a plethora of subsequent methods<sup>[47][48][49][50]</sup>. While most introduce novel architectures or loss functions,<sup>[51]</sup> reformulate depth estimation as an ordinal regression problem through discretization of the depth ranges. Another innovative approach by<sup>[52]</sup> incorporates neural conditional random fields to model contextual dependencies, further refining depth predictions. Besides,<sup>[53]</sup> impose geometric constraints based on piecewise planarity priors.

**Vision Transformers.** Naturally, the rise of vision transformers<sup>[54]</sup> has also significantly impacted the field of MDE. These models employ the self-attention mechanism of the transformer to aggregate depth information across a more extensive field of view to capture long-range dependencies and global context, leading to more accurate and consistent depth maps. Inherently, multiple approaches have successfully adapted vision transformers to MDE<sup>[55][56][57][58][59][60][61][62][63][64][65]</sup>.

**Hybrid Approaches.** Beyond more traditional regression-based methods, some works creatively treat MDE as a combined regression-classification task. By discretizing the depth range into bins, these methods simplify the learning task improve performance in some cases. Notable examples include AdaBins<sup>[66]</sup>, BinsFormer<sup>[67]</sup>, and LocalBins<sup>[10]</sup>.

**Generative Models.** A more recent trend in MDE includes repurposing generative models such as diffusion models<sup>[68][69][70][71][65][72]</sup>, effectively building on its predecessor, generative adversarial networks<sup>[73][74]</sup>.

**Depth in the Wild.** Estimating depth "in the wild" has become an increasingly important area of research in MDE. It refers to the challenge of predicting accurate depth estimates in unconstrained environments, where lighting, scene structure, and camera parameters vary significantly. With the increasing availability of compute resources, researchers have found scaling to be a valuable tool to tackle this challenge. By constructing large and diverse depth datasets and leveraging powerful

foundation models, MDE has become more accessible and robust to real-world use. Foundation models are large neural networks pre-trained on internet-scale data, which allow them to develop a high-level visual understanding that can either be used directly or fine-tuned for a variety of downstream tasks<sup>[75]</sup>.

**Ordinal Depth.** One of the earlier works aimed at addressing depth in the wild by leveraging the scale of data is<sup>[14]</sup>. Building on this idea,<sup>[15]</sup> introduced the OASIS dataset, a large-scale dataset specifically designed for depth and normal estimation. It is worth noting, however, both of these works primarily focus on relative (ordinal) depth, which only estimates the depth order instead of providing absolute measurements. While ordinal depth can provide valuable information about the structure of a given scene, its practical use is limited.

**Affine-invariant Depth.** To overcome the limitations of ordinal depth, several studies have explored affine-invariant depth estimation, which provides depth estimates up to an unknown affine transformation. In other words, the absolute scale and offset of the depth map can vary while the relative depth differences are preserved. For instance, models trained on the MegaDepth dataset<sup>[16]</sup>, which uses multi-view internet photo collections along with structure-from-motion and multi-view stereo methods to create depth maps, can generalize well to unseen images. Another significant contribution is MiDaS<sup>[17]</sup>, which achieves the ability to generalize across a variety of scenes and conditions through training on a mixture of multiple datasets.

**Metric Depth.** For applications that often require absolute distances, such as robotics<sup>[3][4]</sup>, augmented reality<sup>[18]</sup>, and autonomous driving<sup>[1][2]</sup>, metric depth estimation is crucial. Metric depth estimation aims to provide absolute depth measurements in real-world units (e.g., meters or centimeters). Inconveniently, zero-shot generalization is particularly challenging due to the metric ambiguities introduced by different camera models. Aside from some works that explicitly incorporate camera intrinsics as an additional input<sup>[76][77]</sup> to directly solve this issue, current state-of-the-art metric depth estimation approaches still rely on fine-tuning powerful foundation models in the respective domain<sup>[10][11][12][13]</sup>.

## 2.2. Uncertainty Quantification

**Overview.** A wide variety of UQ methods have been proposed to address the shortcomings of deep neural networks, particularly in terms of reliability and robustness<sup>[78][40][42][41][79][80][35][81]</sup>. Among them, sampling-based approaches are the most prominent due to their ease of use and effectiveness in

providing high-quality uncertainty estimates<sup>[78][40][42]</sup>. While they usually produce the most accurate uncertainty estimates, the computational cost associated with the necessity of multiple forward passes often makes them unusable for real-world applications that require fast inference times or that are running on resource-constrained devices.

**Sampling-based Methods.** One of the simplest and most prevalent sampling-based UQ methods is Monte Carlo Dropout (MCD)<sup>[40]</sup>. MCD approximates a Gaussian process by keeping dropout layers active during both training and testing. Originally, dropout layers were solely introduced as a regularization technique to prevent overfitting<sup>[82]</sup>. With MCD, however, they are also used during test time to turn a deterministic model into a stochastic one to sample from the posterior distribution. The predictive uncertainty of a given model can then easily be estimated by calculating the standard deviation (or variance) of the samples. Another widespread sampling-based UQ method are Deep Ensembles<sup>[42]</sup>, which are generally considered the state-of-the-art for UQ across various tasks<sup>[83][84][85][86]</sup>. They consist of a collection of independently trained models, ideally each initialized with random weights and optimized with random data augmentations to maximize the diversity among the ensemble members<sup>[87]</sup>. The high effectiveness of Deep Ensembles comes at the cost of a very high computational overhead due to the need to train and evaluate multiple models.

**Efficiency-oriented Methods.** In addition to these methods, there are several approaches that aim to balance computational efficiency with the quality of uncertainty estimates<sup>[41][79][80][35]</sup>. For instance, Sub-Ensembles<sup>[41]</sup> exploit subnetworks within a single model to produce diverse predictions without the need of a full ensemble of models. They offer an effective trade-off between uncertainty quality and computational cost, which can easily be tuned based on the given constraints.

**UQ in Self-supervised MDE.** While substantial progress has been made in developing effective UQ methods, integrating these techniques with MDE comes with some unique challenges like the limited ground truth data, which is expensive and difficult to acquire, especially at scale. For that reason, a significant body of research has focused on UQ in self-supervised MDE<sup>[88][89][90][91][92]</sup>. While they all explore different strategies to estimate the uncertainty, some even manage to leverage the uncertainty to enhance model performance<sup>[88][90][91]</sup>.

**UQ in Supervised MDE.** For supervised learning scenarios, one common approach to UQ involves modeling the regression output as a parametric distribution and training the model to estimate its parameters<sup>[93][39]</sup>. This approach allows the model to not only output the depth estimate but also to

measure the corresponding uncertainty, typically represented as the variance of the distribution. Similarly,<sup>[94]</sup> propose an auxiliary network that exploits the output and the intermediate representations of the main model to estimate the uncertainty.<sup>[95]</sup> introduce a post hoc uncertainty estimation method that relies on gradients extracted with an auxiliary loss function. This technique utilizes Test-Time Augmentation<sup>[43]</sup> to investigate the correspondence of the depth prediction for an image and its horizontally flipped counterpart. Another innovative approach is proposed by<sup>[96]</sup>, who address computational efficiency by optimizing a set of latent prototypes. The uncertainty is quantified by examining the position of an input sample in the prototype space. Lastly,<sup>[97]</sup> developed a training-free uncertainty estimation approach based on tolerable perturbations during inference and using the variance of multiple outputs as a surrogate for the uncertainty.

**Research Gap.** Despite significant advancements in UQ for MDE, integrating these techniques with large-scale foundation models remains unexplored. We aim to address this gap by combining multiple UQ methods with the state-of-the-art DepthAnythingV2 foundation model<sup>[13]</sup>, enabling pixel-wise uncertainty estimates in addition to metric depth measurements. Our findings will contribute to a more nuanced understanding of the capabilities and limitations of these models, emphasizing the importance of not only striving for higher performance but also making MDE more reliable and trustworthy for real-world use.

### 3. DepthAnything Foundation Model

DepthAnythingV2<sup>[13]</sup> is one the most recent state-of-the-art foundation models for MDE, which can easily be fine-tuned for metric depth estimation, making it the perfect candidate for exploring the combination of various UQ methods with MDE foundation models. Consequently, we want to provide a more detailed overview of this model, which was built upon the framework established by its predecessor, DepthAnythingV1<sup>[12]</sup>.

**DepthAnythingV1.** <sup>[13]</sup> laid the groundwork for creating a versatile foundation model for MDE using the DINOv2 encoder<sup>[98]</sup> for feature extraction with the DPT decoder<sup>[61]</sup> for depth regression. The training process of DepthAnythingV1 involves a semi-supervised approach using a student-teacher framework. The teacher model generates pseudo-labels for an extremely large corpus of unlabeled images (approx. 62 million from 8 public datasets), while the student is trained on both the pseudo-labels and a set of 1.5 million labeled images from 6 public datasets. To ensure robustness of the learned representations,



they additionally apply strong image perturbations for the student. These include strong color distortion like color jittering and Gaussian blurring and strong spatial distortion through CutMix<sup>[99]</sup>.

To further refine the model's capabilities, they also introduce an auxiliary feature alignment loss. It measures the cosine similarity between the features of the student model and those of a frozen DINOv2 encoder, which is a powerful model for semantic-related tasks like image retrieval and semantic segmentation. This potent addition to the training process helps imbue the DepthAnythingV1 model with high-level semantic understanding to further improve the depth estimation.

**DepthAnythingV2.** Building on the success of DepthAnythingV1,<sup>[13]</sup> quickly introduced several key advancements that enable finer and more robust depth predictions. These improvements are centered around the following three strategies:

- 1. Synthetic Data for Label Accuracy:** One of the most significant changes for DepthAnythingV2 is the replacement of all labeled real images with synthetic images. This alteration is motivated by the desire to eliminate label noise and address the lack of detail often ignored in real datasets. In contrast to real images, synthetic data allows for precise depth training and avoids the inconsistencies found in real-world labels.
- 2. Scaling up the Teacher Model:** To mitigate the drawbacks of synthetic images, such as distribution shifts and restricted scene coverage, the capacity of the teacher model significantly increased. DepthAnythingV2 employs DINOv2-G, the most powerful variant of the DINOv2 encoder<sup>[61]</sup>.
- 3. Leveraging Pseudo-Labeled Real Images:** To bridge the gap between synthetic images and the complexity of real-world scenes, DepthAnythingV2 incorporates large-scale pseudo-labeled real images into its training pipeline. This does not only expand the lacking scene coverage of the synthetic images but also ensures that the model is exposed to a wide variety of real-world scenarios, improving its generalization capabilities.

These advancements make DepthAnythingV2 one of the most powerful, if not the most powerful currently available MDE foundation model. Aside from providing high-quality depth predictions across a diverse range of applications, the model can easily be fine-tuned for metric depth estimation tasks, achieving state-of-the-art results. Given these attributes, DepthAnythingV2 serves as an ideal candidate for exploring foundation model uncertainty in MDE. We aim to assess various UQ methods in

combination with this foundation model, enabling pixel-wise uncertainty measures while maintaining high accuracy in the regressing metric depths in real-world applications.

## 4. Methodology

### 4.1. Overview

Our primary research question is straightforward: How can we bridge the gap between ground-breaking results in research and safe deployment in real-world applications that need robust metric depth estimates with corresponding uncertainties.

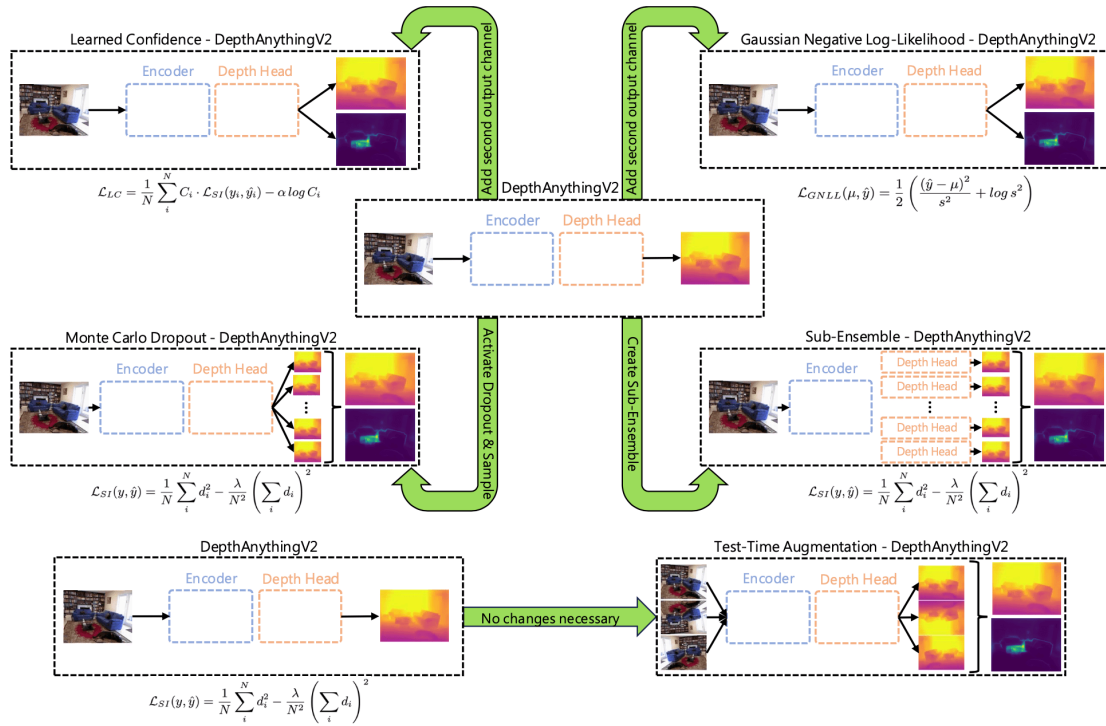
As Figure 2 shows, we study five different approaches to not only estimate metric depths but also their corresponding uncertainties: Learned Confidence (LC)<sup>[38]</sup>, Gaussian Negative Log-Likelihood (GNLL)<sup>[39]</sup>, Monte Carlo Dropout (MCD)<sup>[40]</sup>, Sub-Ensembles (SE)<sup>[41]</sup>, and Test-Time Augmentation (TTA)<sup>[43]</sup>

The upper two approaches of Figure 2, **LC** and **GNLL**, are fairly simple since they only require adding a second output channel to the already existing depth head. The first output channel outputs the metric depth maps and the second output channel the uncertainty.

The third option, **MCD**, is equally straightforward since it only requires activating all of the already existing dropout layers in the model while fine-tuning. During inference, these dropout layers are kept active and multiple depth outputs are being sampled. By computing the mean and variance, the final depth map and the corresponding uncertainty are obtained.

The fourth option, **SE**, is possibly the most complicated and requires significant changes to the architecture. Instead of using just one depth head, a sub-ensemble of randomly initialized depth heads is created. During inference, every depth head predicts slightly different depth samples. Similar to the third option, **MCD**, the final depth map and uncertainty can be obtained by computing the mean and variance.

Finally, as Figure 2 shows, we also examine **TTA**, which does not require any changes to the fine-tuning process of DepthAnythingV2. Instead, we apply horizontal and vertical flipping during test-time to create two additional inputs to create a total of three unique depth samples. Based on these, we compute the mean and variance to obtain the final depth map and the uncertainty.



**Figure 2.** A schematic overview of how to fuse the five different uncertainty quantification approaches with the DepthAnythingV2<sup>[13]</sup> foundation model.

#### 4.2. Learned Confidence

The general approach of LC was originally proposed by<sup>[38]</sup> for classification tasks, but has already been adapted before for regression tasks by<sup>[100]</sup>. The confidences, which we interpret as uncertainties, can simply be learned in addition to the primary objective function:

$$\mathcal{L}_{LC} = \frac{1}{N} \sum_i^N C_i \cdot \mathcal{L}_{SI}(y_i, \hat{y}_i) - \alpha \log C_i, \quad (1)$$

where  $N$  represents the number of pixels having valid ground truth values,  $C_i$  denotes the confidence score of the  $i$ -th valid pixel generated by the model,  $y_i$  is the depth output of the model,  $\hat{y}_i$  is the ground truth depth,  $\alpha$  is set to 0.2, following<sup>[100]</sup>, and  $\mathcal{L}_{SI}(y_i, \hat{y}_i)$  is the scale-invariant loss introduced by<sup>[46]</sup>:

$$\mathcal{L}_{SI}(y, \hat{y}) = \frac{1}{N} \sum_i^N d_i^2 - \frac{\lambda}{N^2} \left( \sum_i^N d_i \right)^2, \quad (2)$$

where  $d_i = \log y_i - \log \hat{y}_i$ . We follow the fine-tuning recommendations of DepthAnythingV2<sup>[13]</sup> and set  $\lambda = 0.15$ .

### 4.3. Gaussian Negative Log-Likelihood Loss

For depth regression, neural networks are usually only trained to output a predictive mean  $\mu$ . To also approximate the corresponding variance  $s^2$ , i.e. the uncertainty, we follow the approach of<sup>[39]</sup>. By treating the neural network prediction as a sample from a Gaussian distribution, we can minimize the Gaussian Negative Log-Likelihood (GNLL) loss, which can be formulated as:

$$\mathcal{L}_{GNLL}(\mu, \hat{y}) = \frac{1}{2} \left( \frac{(\hat{y} - \mu)^2}{s^2} + \log s^2 \right). \quad (3)$$

Analogous to Equation 1 for LC, there is no ground truth for the uncertainty, which means that  $s^2$  is solely learned implicitly through the optimization of the predictive means  $\mu$  based on the ground truth labels  $\hat{y}$ .

### 4.4. MC Dropout

Using MCD<sup>[40]</sup> to estimate the predictive mean  $\mu$  and the corresponding uncertainty, i.e., the variance  $s^2$  or standard deviation  $s$ , is fairly straightforward. Since the DepthAnythingV2 model already applies dropout layers throughout its architecture, we simply have to activate them not only during training but also during inference to sample from the posterior of the network.

To compute the predictive mean  $\mu$ , we take the average of all the samples:

$$\mu = \frac{1}{T} \sum_t^T y_t, \quad (4)$$

where  $T$  is the number of samples and  $y_t$  is the  $t$ -th depth prediction of the network.

For the uncertainty, we calculate the variance:

$$s^2 = \frac{1}{T-1} \sum_t^T (y_t - \mu)^2. \quad (5)$$

Besides that, we follow the fine-tuning recommendations of DepthAnythingV2<sup>[13]</sup>, using the scale-invariant loss  $\mathcal{L}_{SI}(y, \hat{y})$  from Equation 2 as the objective function.

## 4.5. Sub-Ensemble

Sub-Ensembles (SE)<sup>[41]</sup> enable the approximation of Deep Ensemble. While a Deep Ensemble requires multiple models to be trained and used during inference, the SE only requires a subset of the layers to be multiplied. As shown by Figure 2, we use a shared encoder for multiple randomly initialized depth heads. To maximize the diversity across the depth heads and decrease the training time, we cycle through the heads during training. Per training batch, only one head is optimized, the others are ignored. During inference, however, each depth head predicts a unique sample  $y_t$  based on the extracted feature from the encoder. Similar to MCD, we can compute the mean  $\mu$  and variance  $s^2$  (see Equations 4 and 5) to get the desired output, i.e., a final depth map and a corresponding uncertainty.

As for MCD, we follow the fine-tuning recommendations of DepthAnythingV2<sup>[13]</sup>, using the scale-invariant loss  $\mathcal{L}_{SI}(y, \hat{y})$  from Equation 2 as the objective function.

## 4.6. Test-Time Augmentation

In contrast to the other four uncertainty quantification approaches, we just fine-tune the DepthAnythingV2 model with the scale-invariant loss  $\mathcal{L}_{SI}$  from Equation 2 for metric depth estimation and apply Test-Time Augmentation (TTA) after training<sup>[43]</sup>. As shown by Figure 2, we flip the input image vertically as well as horizontally and perform inference with each. As a result, we obtain three unique depth samples  $y_t$  that we can use to compute the the mean  $\mu$  and variance  $s^2$  (see Equations 4 and 5).

# 5. Experiments

## 5.1. Experimental Setup

**Training.** For all training processes, we follow the default settings of DepthAnythingV2 for metric depth fine-tuning<sup>[13]</sup>, using an AdamW optimizer<sup>[101]</sup> with a base learning rate of  $6 \cdot 10^{-5}$ , a weight decay of 0.01, and a polynomial learning rate scheduler:

$$\text{lr} = \text{lr}_{\text{base}} \cdot \left( 1 - \frac{\text{iteration}}{\text{total iterations}} \right)^{0.9}, \quad (6)$$

where lr is the current learning rate and  $\text{lr}_{\text{base}}$  is the initial base learning rate. Every model is trained for 25 epochs with an effective batch size of 16, using four NVIDIA A100 GPUs. We do not employ any early stopping techniques and hence only evaluate the final model checkpoints.

Scene / Application	Resolution (W × H)	Training Images	Test Images	
Cityscapes <sup>[21]</sup>	Outdoor	2048 × 1024	2975	500
NYUv2 <sup>[19]</sup>	Indoor	640 × 480	795	654
UseGeo <sup>[45]</sup>	Aerial	1989 × 1320	551	277
HOPE <sup>[44]</sup>	Robotics	1920 × 1080	49450 (synth.)	457 (real)

**Table 1.** Overview of metric depth datasets that we used for evaluation.

**Datasets.** We conduct our experiments on four highly different datasets, simulating a broad range of real-world applications, as shown by Table 1. Cityscapes<sup>[21]</sup> provides an urban street scene benchmark dataset with high-resolution images. In contrast, NYUv2<sup>[19]</sup> presents indoor scenes with a very low image resolution. UseGeo<sup>[45]</sup> covers high-resolution aerial images, which are often neglected in the computer vision community despite their significance in many real-world applications. Finally, the HOPE<sup>[44]</sup> dataset offers a variety of household objects, originally designed for pose estimation. The main reason why the HOPE dataset is so interesting is that the training dataset is based on almost 50,000 synthetic images, whereas the test dataset consists of just 457 real images. A unique challenge that is often overlooked but fairly common, especially in robotics<sup>[102][103][104]</sup>.

**Data Augmentations.** Regardless of the trained model, we apply random cropping with a crop size of  $756 \times 756$  pixels on all datasets except NYUv2, which uses  $630 \times 476$  pixels, and random horizontal flipping with a flip chance of 50%. For testing purposes, we use the original image resolutions as shown by Table 1.

**Metrics.** For quantitative evaluations of the metric depth estimation, we report all common metrics: root mean squared error (RMSE), absolute relative error (AbsRel), logarithmic root mean squared error (log10), and three threshold-based accuracies ( $\delta_1, \delta_2, \delta_3$ ).

To evaluate the uncertainty quality, we exploit the following uncertainty evaluation metrics proposed by<sup>[34]</sup>:

1. **p(accurate|certain)**: The probability that the model is accurate on its output given that the uncertainty is below a specified threshold.
2. **p(uncertain|inaccurate)**: The probability that the uncertainty of the model exceeds a specified threshold given that the prediction is inaccurate.
3. **PAvPU**: The combination of metrics 1 and 2.

Despite the fact that these metrics have originally been proposed for semantic segmentation<sup>[34]</sup>, they can also be used to evaluate the depth regression uncertainties<sup>[105][86]</sup>. To determine whether a depth prediction is accurate or inaccurate, we use the strictest threshold-based accuracy  $\delta_1$ :

$$\max\left(\frac{y}{\hat{y}}, \frac{\hat{y}}{y}\right) = \delta_1 < 1.25. \quad (7)$$

To simulate real-world employment, we set the threshold to determine whether a pixel is certain or uncertain to the median uncertainty of a given image<sup>[86]</sup>.

**Uncertainty Quantification.** Unless stated otherwise, for MCD and SE, the final depth map and corresponding uncertainty are computed using ten samples or ten depth heads, respectively, in accordance with the findings in<sup>[42][106][31][36]</sup>.

## 5.2. Quantitative Evaluation

NYUv2 (indoor)		Trainable Parameters [M] ↓	FLOPs [G] ↓	Training Time [mm:ss] ↓	Inference Time [ms] ↓	FPS ↑
ViT-S	Baseline	24.8	66.82	00:27	12.9 ± 0.1	77.5
	TTA	24.8	66.82	00:27	40.0 ± 0.9	25.0
	LC	24.8	66.83	00:27	12.9 ± 0.2	77.5
	GNLL	24.8	66.83	00:27	12.9 ± 0.2	77.5
	SE	49.3	177.29	00:29	36.5 ± 0.5	27.4
	MCD	24.8	66.82	00:30	128.4 ± 0.5	7.8
ViT-B	Baseline	97.5	217.49	00:48	24.9 ± 0.5	40.2
	TTA	97.5	217.49	00:48	75.5 ± 1.2	13.2
	LC	97.5	217.50	00:49	25.1 ± 0.5	39.8
	GNLL	97.5	217.50	00:49	25.0 ± 0.7	40.0
	SE	195.5	608.09	00:52	61.9 ± 0.8	16.2
	MCD	97.5	217.49	00:54	248.4 ± 2.7	4.0
ViT-L	Baseline	335.3	741.40	02:11	57.1 ± 3.5	17.5
	TTA	335.3	741.40	02:11	172.6 ± 11.6	5.8
	LC	335.3	741.41	02:14	57.0 ± 3.9	17.5
	GNLL	335.3	741.41	02:13	57.1 ± 3.3	17.5
	SE	613.8	2204.07	02:17	131.6 ± 5.2	7.6
	MCD	335.3	741.40	02:26	569.8 ± 24.3	1.8

**Table 2.** Efficiency comparison between the five chosen uncertainty quantification methods: Test-Time Augmentation (TTA), Learned Confidence (LC), Gaussian Negative Log-Likelihood (GNLL), Sub-Ensemble (SE), and MC Dropout (MCD) for three different encoder sizes: ViT-S, ViT-B, ViT-L on the NYUv2<sup>[19]</sup> dataset. We compare the number of trainable parameters, FLOPs, training time per epoch on a



single A100 GPU, inference time per image, and the respective FPS. The mean inference time and corresponding standard deviation are based on 1000 forward passes.

**Efficiency.** As shown by Table 2, training times are comparable to the baseline for all methods, with LC and GNLL matching it exactly in both training and inference time. SE and MCD increase training times by around 5% - 10%. While SE nearly triples inference time and roughly doubles the trainable parameters, MCD requires roughly 10 times the inference time due to the costly sampling process that roughly scales linearly with the amount of samples. TTA also triples inference time as it requires three forward passes. Overall, LC and GNLL are the most efficient approaches as they do not require any additional computational overhead compared to the baseline. These experiments were conducted on the NYUv2 dataset only, so exact numbers may vary across datasets, but the general findings should be representative.

NYUv2 (indoor)		RMSE ↓	AbsRel ↓	log10 ↓	$\delta_1$ ↑	$\delta_2$ ↑	$\delta_3$ ↑	p(acc cer) ↑	p(unc ina) ↑	PAvPU ↑
ViT-S	Baseline	<b>0.340</b>	0.093	0.039	0.928	0.988	0.997	-	-	-
	TTA	0.399	0.111	0.049	0.881	0.984	0.997	0.903	0.602	0.522
	LC	0.343	0.090	0.039	0.930	0.988	0.997	0.920	0.369	0.490
	GNLL	0.342	0.094	0.040	0.924	0.987	0.997	<b>0.953</b>	<b>0.846</b>	0.529
	SE	<b>0.340</b>	0.092	0.039	0.926	0.988	0.997	0.937	0.704	0.511
	MCD (10%)	0.422	0.121	0.050	0.867	0.973	0.992	0.900	0.692	<b>0.533</b>
ViT-B	Baseline	0.307	0.080	0.034	0.948	0.991	0.998	-	-	-
	TTA	0.359	0.099	0.435	0.910	0.988	0.998	0.924	0.599	0.514
	LC	0.314	0.085	0.036	0.943	0.991	0.998	0.926	0.292	0.483
	GNLL	<b>0.305</b>	0.079	0.034	0.949	0.991	0.998	<b>0.966</b>	<b>0.826</b>	0.517
	SE	0.309	0.080	0.034	0.947	0.991	0.998	0.955	0.698	0.507
	MCD (10%)	0.339	0.091	0.039	0.925	0.986	0.997	0.952	0.774	<b>0.527</b>
ViT-L	Baseline	<b>0.270</b>	0.068	0.030	0.964	0.993	0.998	-	-	-
	TTA	0.324	0.087	0.039	0.938	0.992	0.998	0.944	0.598	0.507
	LC	0.275	0.069	0.030	0.963	0.993	0.998	0.946	0.268	0.483
	GNLL	0.285	0.072	0.031	0.959	0.992	0.998	<b>0.980</b>	<b>0.912</b>	0.521
	SE	0.280	0.070	0.030	0.961	0.993	0.998	0.969	0.734	0.508
	MCD (10%)	0.339	0.091	0.039	0.925	0.986	0.997	0.967	0.798	<b>0.522</b>

**Table 3.** Quantitative comparison on the NYUv2<sup>[19]</sup> dataset between the five chosen uncertainty quantification methods: Test-Time Augmentation (TTA), Learned Confidence (LC), Gaussian Negative Log-Likelihood (GNLL), Sub-Ensemble (SE), and MC Dropout (MCD) for three different encoder sizes: ViT-S, ViT-B, ViT-L. Best results for RMSE and the three uncertainty metrics are marked in **bold** for each encoder.

NYUv2. Table 3 shows a quantitative comparison for the NYUv2 dataset<sup>[19]</sup>. LC, GNLL, and SE maintain depth quality similar to the baseline, with SE and GNLL even surpassing it for the ViT-S and ViT-B encoders, respectively. TTA and MCD exhibit a slight degradation but remain competitive.

Regarding the uncertainty quality, GNLL emerges as the top-performing method, consistently outperforms all others in terms of  $p(\text{acc|cer})$  and  $p(\text{unc|ina})$  across all three encoder sizes, achieving impressive values of up to 98.0% and 91.2%, respectively. For PAVPU, both GNLL and MCD deliver the best results. While all other methods are somewhat competitive with each other, LC clearly falls behind with regard to  $p(\text{unc|ina})$ , achieving only 26.8% for ViT-L.

Cityscapes (outdoor)		RMSE ↓	AbsRel ↓	log10 ↓	$\delta_1$ ↑	$\delta_2$ ↑	$\delta_3$ ↑	p(acc cer) ↑	p(unc ina) ↑	PAvPU ↑
ViT-S	Baseline	7.138	0.219	0.084	0.704	0.964	0.991	-	-	-
	TTA	<b>6.474</b>	0.223	0.089	0.576	0.965	0.991	0.335	0.421	0.410
	LC	7.492	0.206	0.078	0.733	0.958	0.991	0.658	0.619	0.611
	GNLL	7.739	0.236	0.089	0.669	0.949	0.987	<b>0.692</b>	<b>0.706</b>	<b>0.695</b>
	SE	7.714	0.234	0.088	0.681	0.954	0.990	0.658	0.654	0.649
	MCD (10%)	8.527	0.307	0.113	0.417	0.920	0.984	0.328	0.525	0.516
ViT-B	Baseline	6.884	0.252	0.095	0.599	0.965	0.992	-	-	-
	TTA	<b>5.757</b>	0.247	0.095	0.526	0.967	0.993	0.288	0.420	0.396
	LC	7.711	0.285	0.107	0.434	0.958	0.991	0.329	0.506	0.502
	GNLL	7.092	0.244	0.094	0.613	0.966	0.991	<b>0.594</b>	<b>0.635</b>	<b>0.629</b>
	SE	7.824	0.271	0.102	0.534	0.957	0.991	0.490	0.587	0.588
	MCD (10%)	8.268	0.288	0.107	0.488	0.939	0.989	0.449	0.579	0.583
ViT-L	Baseline	6.655	0.256	0.097	0.558	0.972	0.993	-	-	-
	TTA	<b>5.298</b>	0.227	0.088	0.608	0.979	0.994	0.371	0.431	0.416
	LC	7.392	0.280	0.105	0.416	0.969	0.993	0.292	0.488	0.488
	GNLL	6.562	0.234	0.092	0.628	0.970	0.990	<b>0.581</b>	<b>0.620</b>	<b>0.607</b>
	SE	7.522	0.272	0.103	0.500	0.966	0.993	0.446	0.568	0.570
	MCD (10%)	8.268	0.288	0.107	0.488	0.939	0.989	0.480	0.584	0.584

**Table 4.** Quantitative comparison on the Cityscapes<sup>[21]</sup> dataset between the five chosen uncertainty quantification methods: Test-Time Augmentation (TTA), Learned Confidence (LC), Gaussian Negative Log-Likelihood (GNLL), Sub-Ensemble (SE), and MC Dropout (MCD) for three different encoder sizes: ViT-S, ViT-B, ViT-L. Best results for RMSE and the three uncertainty metrics are marked in **bold** for each encoder.

**Cityscapes.** For the Cityscapes dataset<sup>[21]</sup>, as shown by Table 4, TTA stands out by significantly outperforming the baseline across all three encoder sizes in terms of depth quality. In contrast, LC, GNLL, and SE exhibit noticeable degradation, while MCD performs the worst.

For uncertainty quality, GNLL is the standout performer, decisively surpassing all other methods for all three metrics and encoder sizes. It achieves remarkable values of 69.2% for  $p(\text{acc}|\text{cer})$ , 70.6% for  $p(\text{unc}|\text{ina})$ , and 69.5% for PAvPU, setting a benchmark for reliability in this dataset. The remaining methods deliver less consistent results, with TTA generally performing the worst, showing values as low as 28.8% for  $p(\text{acc}|\text{cer})$ , 42.0% for  $p(\text{unc}|\text{ina})$ , and 39.6% for PAvPU. Interestingly, GNLL remains resilient despite the observed uncertainty degradation with increasing encoder size for most other methods.

UseGeo (aerial)		RMSE ↓	AbsRel ↓	log10 ↓	$\delta_1$ ↑	$\delta_2$ ↑	$\delta_3$ ↑	p(acc cer) ↑	p(unc ina) ↑	PAvPU ↑
ViT-S	Baseline	7.366	0.077	0.032	0.973	0.994	0.999	-	-	-
	TTA	7.078	0.074	0.031	0.977	0.995	0.999	0.976	0.647	0.500
	LC	8.213	0.086	0.036	0.958	0.995	1.000	0.963	0.481	<b>0.506</b>
	GNLL	7.467	0.076	0.032	0.979	0.993	0.998	0.976	0.237	0.500
	SE	7.259	0.076	0.032	0.976	0.994	0.999	0.971	0.467	0.495
	MCD (10%)	<b>6.682</b>	0.068	0.029	0.982	0.994	0.998	<b>0.982</b>	<b>0.672</b>	0.501
ViT-B	Baseline	6.386	0.067	0.028	0.981	0.995	0.999	-	-	-
	TTA	<b>6.060</b>	0.631	0.027	0.985	0.995	0.999	<b>0.984</b>	0.609	0.499
	LC	6.612	0.068	0.029	0.975	0.995	1.000	0.967	0.454	0.492
	GNLL	7.810	0.080	0.035	0.972	0.990	0.997	0.971	0.293	0.499
	SE	6.491	0.068	0.028	0.980	0.994	0.999	0.981	0.559	0.502
	MCD (10%)	6.641	0.070	0.029	0.978	0.994	0.999	0.982	<b>0.657</b>	<b>0.504</b>
ViT-L	Baseline	6.173	0.065	0.027	0.981	0.995	0.999	-	-	-
	TTA	5.898	0.063	0.026	0.982	0.995	0.999	0.980	0.596	0.499
	LC	<b>5.406</b>	0.056	0.024	0.986	0.995	1.000	<b>0.985</b>	0.477	0.500
	GNLL	7.082	0.073	0.031	0.980	0.991	0.998	0.980	0.294	0.498
	SE	6.260	0.067	0.028	0.980	0.995	1.000	0.977	0.577	0.497
	MCD (10%)	6.697	0.071	0.029	0.981	0.995	0.999	0.982	<b>0.669</b>	<b>0.501</b>

**Table 5.** Quantitative comparison on the UseGeo<sup>[45]</sup> dataset between the five chosen uncertainty quantification methods: Test-Time Augmentation (TTA), Learned Confidence (LC), Gaussian Negative Log-Likelihood (GNLL), Sub-Ensemble (SE), and MC Dropout (MCD) for three different encoder sizes: ViT-S, ViT-B, ViT-L. Best results for RMSE and the three uncertainty metrics are marked in **bold** for each encoder.

UseGeo. For the UseGeo dataset<sup>[45]</sup>, as presented by Table 5, the depth quality results are inconsistent, with methods sometimes surpassing and at other times falling short of the baseline. TTA is the only approach that consistently outperforms the baseline across all three encoder sizes.

In terms of uncertainty quality, all methods deliver near-perfect results for  $p(\text{acc|cer})$  of at least 96.3%. For  $p(\text{unc|ina})$  and PAVPU, MCD generally performs best with values of up 67.2% and 50.4%, respectively. In contrast to the other datasets, GNLL significantly lags behind the other methods on UseGeo. This is likely due to the large depth values in UseGeo, which led to much higher absolute GNLL loss values. The GNLL loss incorporates a logarithmic term that penalizes high uncertainty estimates, and with large depth values, the uncertainties—and hence the loss—are naturally magnified. Notably, no hyperparameter adjustments were made to address this, ensuring comparability but potentially hindering GNLL’s optimization in this particular case.

HOPE (robotics)		RMSE ↓	AbsRel ↓	log10 ↓	$\delta_1$ ↑	$\delta_2$ ↑	$\delta_3$ ↑	p(acc cer) ↑	p(unc ina) ↑	PAvPU ↑
ViT-S	Baseline	0.263	0.265	0.115	0.537	0.821	0.942	-	-	-
	TTA	0.264	0.262	0.114	0.539	0.818	0.943	<b>0.421</b>	0.564	0.552
	LC	<b>0.259</b>	0.262	0.114	0.537	0.822	0.945	0.339	0.476	0.474
	GNLL	0.277	0.287	0.124	0.492	0.795	0.929	0.404	<b>0.575</b>	<b>0.567</b>
	SE	0.262	0.263	0.117	0.544	0.809	0.933	0.393	0.528	0.522
	MCD (10%)	0.274	0.286	0.118	0.514	0.816	0.937	0.398	0.562	0.553
ViT-B	Baseline	0.222	0.232	0.094	0.616	0.892	0.969	-	-	-
	TTA	0.221	0.230	0.093	0.621	0.891	0.968	0.462	0.573	0.558
	LC	0.224	0.227	0.095	0.616	0.883	0.966	0.330	0.419	0.427
	GNLL	0.223	0.225	0.094	0.619	0.892	0.972	<b>0.497</b>	<b>0.622</b>	<b>0.596</b>
	SE	<b>0.218</b>	0.230	0.094	0.625	0.889	0.967	0.448	0.546	0.543
	MCD (10%)	0.245	0.269	0.106	0.560	0.851	0.955	0.405	0.555	0.543
ViT-L	Baseline	0.223	0.238	0.096	0.588	0.906	0.980	-	-	-
	TTA	0.217	0.232	0.094	0.599	0.904	0.980	0.436	0.576	0.557
	LC	<b>0.215</b>	0.226	0.092	0.604	0.911	0.980	0.325	0.426	0.441
	GNLL	0.229	0.244	0.099	0.588	0.888	0.973	<b>0.460</b>	<b>0.608</b>	<b>0.586</b>
	SE	0.235	0.252	0.098	0.594	0.893	0.974	0.442	0.574	0.569
	MCD (10%)	0.249	0.272	0.107	0.557	0.859	0.966	0.416	0.580	0.563

**Table 6.** Quantitative comparison on the HOPE<sup>[44]</sup> dataset between the five chosen uncertainty quantification methods: Test-Time Augmentation (TTA), Learned Confidence (LC), Gaussian Negative Log-Likelihood (GNLL), Sub-Ensemble (SE), and MC Dropout (MCD) for three different encoder sizes: ViT-S, ViT-B, ViT-L. Best results for RMSE and the three uncertainty metrics are marked in **bold** for each encoder.




Table 6 shows the final quantitative comparison for the HOPE dataset<sup>[44]</sup>. In terms of depth quality, there are only marginal differences between the baseline and all uncertainty approaches, with RMSE values ranging between 0.215 to 0.277.

Regarding the uncertainty quality, GNLL once again asserts itself as the best option for all three metrics across all three encoder sizes, with only one minor exception. GNLL achieves results of up to 49.7% for  $p(\text{acc}|\text{cer})$ , 62.2% for  $p(\text{unc}|\text{ina})$ , and 59.6% for PAVPU. Mirroring its strong performance on NYUv2 (cf. Table 3), GNLL's dominance is evident, while the other methods remain fairly competitive with each other. LC, however, lags significantly behind across all three uncertainty quality metrics, underscoring GNLL's reliability and consistency.

### *5.3. Qualitative Evaluation*

We provide handpicked qualitative examples of four different UQ methods with varying encoder sizes for all four datasets in Figure 3, highlighting the potential for foundation model uncertainty in metric MDE.



**Figure 3.** Qualitative examples for indoor<sup>[19]</sup>, outdoor<sup>[21]</sup>, aerial<sup>[45]</sup>, and robotics<sup>[44]</sup> scenarios with varying uncertainty quantification approaches and encoder sizes. Red rectangles are added to highlight interesting areas. Best viewed in color.

**NYUv2.** The first row displays the GNLL (ViT-S) results, demonstrating high prediction quality overall. Uncertainty is notably elevated around object boundaries and the two open doors in the background, suggesting these regions are likely edge cases. This likely stems from the model’s limited exposure to such depth ranges during training since we limit the maximum depth to just 5m, which is common practice on NYUv2.

Additionally, the second row presents results from TTA (ViT-S), revealing mixed results in terms of prediction and uncertainty quality. While the model assigns high uncertainties to the background, which includes some erroneous predictions, it fails to recognize its substantial prediction error on the backrest of the chair at the bottom of the image. Based on this qualitative comparison, GNLL seems to provide more meaningful uncertainties, corroborating the quantitative findings of Section 5.2.

**Cityscapes.** In the third row, the MCD (ViT-B) predictions exhibit multiple errors, particularly in the top right corner. However, the predicted uncertainties in this area show a correlation between

uncertainty and challenging regions, reinforcing the model's awareness of its limitations.

**UseGeo.** The fourth row shows SE (ViT-B) results, where relative depth predictions are plausible despite reduced accuracy in absolute terms. Across the entire image, especially in the bottom right, uncertainties are heightened for the buildings, emphasizing that the model is aware of key areas where depth errors are most likely.

**HOPE.** The fifth row presents LC (ViT-L) results, where the model struggles with the absolute depth of the large foreground object. At the same time, the entire object is highlighted by high uncertainty, reflecting the model's strong awareness of its own prediction error in this case.

## 6. Conclusion

Motivated by the need to bridge the gap between cutting-edge research and the safe deployment of MDE models in real-world applications, we conducted a comprehensive evaluation of multiple UQ methods in conjunction with the state-of-the-art DepthAnythingV2 foundation model. Our evaluation covered five different UQ approaches – Learned Confidence, Gaussian Negative Log-Likelihood, MC Dropout, Sub-Ensembles, and Test-Time Augmentation – and was carried out across four diverse dataset, covering various domains relevant to real-world applications: NYUv2, Cityscapes, UseGeo, and HOPE.

Our findings highlight fine-tuning with GNLL as the most promising option, consistently delivering high-quality uncertainty estimates while maintaining depth performance comparable to the baseline. Its computational efficiency, which matches that of the baseline, further underscores its practical suitability for deployment.

This study emphasizes the importance and feasibility of integrating UQ into machine vision models, demonstrating that achieving reliable uncertainty estimates need not come at the expense of predictive performance or computational complexity. By addressing this critical aspect, we aim to inspire future research that prioritizes not only performance but also explainability through uncertainty awareness, fostering the development of safer and more reliable models for not only MDE but also other real-world tasks, such as semantic segmentation or pose estimation.

## Acknowledgements

The authors acknowledge support by the state of Baden-Württemberg through bwHPC.

This work is supported by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT partition.

The work is partially supported by the Office of Naval Research (N00014-23-1-2670).

## References

1. <sup>a, b, c</sup>Xue F, Zhuo G, Huang Z, Fu W, Wu Z, Ang MH. Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications. *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE; 2020. p. 2330–2337.
2. <sup>a, b, c</sup>Xiang J, Wang Y, An L, Liu H, Wang Z, Liu J (2022). "Visual attention-based self-supervised absolute depth estimation using geometric priors in autonomous driving". *IEEE Robotics and Automation Letters*. 7 (4): 11998–12005.
3. <sup>a, b, c</sup>Dong X, Garratt MA, Anavatti SG, Abbass HA (2022). "Towards real-time monocular depth estimation for robotics: A survey". *IEEE Transactions on Intelligent Transportation Systems*. 23 (10): 16940–16961.
4. <sup>a, b, c</sup>Roussel T, Van Eycken L, Tuytelaars T. Monocular depth estimation in new environments with absolute scale. *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE; 2019. p. 1735–1741.
5. <sup>^</sup>Zhang L, Rao A, Agrawala M (2023). "Adding conditional control to text-to-image diffusion models". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3836–3847.
6. <sup>^</sup>Liew JH, Yan H, Zhang J, Xu Z, Feng J (2023). "Magicedit: High-fidelity and temporally coherent video editing". *arXiv preprint arXiv:2308.14749*.
7. <sup>^</sup>Xu D, Jiang Y, Wang P, Fan Z, Wang Y, Wang Z (2023). "Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4479--4489.
8. <sup>^</sup>Shahbazi M, Claessens L, Niemeyer M, Collins E, Tonioni A, Van Gool L, Tombari F (2024). "InseRF: Text-Driven Generative Object Insertion in Neural 3D Scenes". *arXiv preprint arXiv:2401.05335*.
9. <sup>^</sup>Shriram J, Trevithick A, Liu L, Ramamoorthi R (2024). "Realmdreamer: Text-driven 3d scene generation with inpainting and depth diffusion". *arXiv preprint arXiv:2404.07199*.

10. <sup>a, b, c, d, e</sup>Bhat SF, Birkel R, Wofk D, Wonka P, Müller M (2023). "Zoedepth: Zero-shot transfer by combining relative and metric depth". *arXiv preprint arXiv:2302.12288*.
11. <sup>a, b, c, d</sup>Piccinelli L, Yang YH, Sakaridis C, Segu M, Li S, Van Gool L, Yu F. UniDepth: Universal Monocular Metric Depth Estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024 June;10106–10116.
12. <sup>a, b, c, d, e</sup>Yang L, Kang B, Huang Z, Xu X, Feng J, Zhao H. Depth anything: Unleashing the power of large-scale unlabeled data. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024;10371–10381.
13. <sup>a, b, c, d, e, f, g, h, i, j, k, l, m, n</sup>Yang L, Kang B, Huang Z, Zhao Z, Xu X, Feng J, Zhao H (2024). "Depth Anything V2". *arXiv preprint arXiv:2406.09414*.
14. <sup>a, b</sup>Chen W, Fu Z, Yang D, Deng J (2016). "Single-image depth perception in the wild". *Advances in neural information processing systems*. 29.
15. <sup>a, b</sup>Chen W, Qian S, Fan D, Kojima N, Hamilton M, Deng J (2020). "Oasis: A large-scale dataset for single image 3d in the wild". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 679–688.
16. <sup>a, b</sup>Li Z, Snavely N (2018). "Megadepth: Learning single-view depth prediction from internet photos". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2041–2050.
17. <sup>a, b</sup>Ranftl R, Lasinger K, Hafner D, Schindler K, Koltun V (2020). "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer". *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 44(3): 1623–1637.
18. <sup>a, b</sup>Kalia M, Navab N, Salcudean T. A real-time interactive augmented reality depth estimation technique for surgical robotics. 2019 *International Conference on Robotics and Automation (ICRA)*. IEEE; 2019. p. 8291–8297.
19. <sup>a, b, c, d, e, f, g, h, i</sup>Silberman N, Hoiem D, Kohli P, Fergus R (2012). "Indoor segmentation and support inference from rgb-d images". *Computer Vision – ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*. Springer, 746–760.
20. <sup>a</sup>Geiger A, Lenz P, Urtasun R. Are we ready for autonomous driving? the kitti vision benchmark suite. 2012 *IEEE conference on computer vision and pattern recognition*. IEEE; 2012. p. 3354–3361.
21. <sup>a, b, c, d, e, f, g</sup>Cordts M, Omran M, Ramos S, Rehfeld T, Enzweiler M, Benenson R, Franke U, Roth S, Schiele B. The cityscapes dataset for semantic urban scene understanding. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016;3213–3223.

22. <sup>△</sup>Song S, Lichtenberg SP, Xiao J. Sun rgb-d: A rgb-d scene understanding benchmark suite. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015:567–576.
23. <sup>△</sup>Roy AG, Conjeti S, Navab N, Wachinger C, Alzheimer's Disease Neuroimaging Initiative, et al. Bayesian QuickNAT: Model uncertainty in deep whole-brain segmentation for structure-wise quality control. *NeuroImage*. 2019; 195:11–22.
24. <sup>△</sup>Gawlikowski J, Tassi CRN, Ali M, Lee J, Humt M, Feng J, Kruspe A, Triebel R, Jung P, Roscher R, Shahzad M, Yang W, Bamler R, Zhu XX (2022). "A Survey of Uncertainty in Deep Neural Networks". arXiv:2107.03342. Available from: <https://arxiv.org/abs/2107.03342>.
25. <sup>△</sup>Lee J, Feng J, Humt M, Müller MG, Triebel R (2022). "Trust your robots! predictive uncertainty estimation of neural networks with sparse gaussian processes". *Conference on Robot Learning, PMLR*, 1168–1179.
26. <sup>△</sup>Lee K, Lee H, Lee K, Shin J (2017). "Training confidence-calibrated classifiers for detecting out-of-distribution samples". arXiv preprint arXiv:1711.09325.
27. <sup>△</sup>Guo C, Pleiss G, Sun Y, Weinberger KQ. On calibration of modern neural networks. *International conference on machine learning. PMLR*; 2017. p. 1321–1330.
28. <sup>△</sup>Rawat M, Wistuba M, Nicolae MI. Harnessing model uncertainty for detecting adversarial examples. *NI PS Workshop on Bayesian Deep Learning*. 2017.
29. <sup>△</sup>Serban AC, Poll E, Visser J (2018). "Adversarial examples—a complete characterisation of the phenomenon". arXiv preprint arXiv:1810.01185. Available from: <https://arxiv.org/abs/1810.01185>.
30. <sup>△</sup>Smith L, Gal Y (2018). "Understanding measures of uncertainty for adversarial example detection". arXiv preprint arXiv:1803.08533.
31. <sup>△</sup><sup>a</sup> Landgraf S, Würsthorn K, Hillemann M, Ulrich M (2024). "Dudes: Deep uncertainty distillation using ensembles for semantic segmentation". *PFG—Journal of Photogrammetry, Remote Sensing and Geoinformation Science*. 92 (2): 101–114.
32. <sup>△</sup>Leibig C, Allken V, Ayhan MS, Berens P, Wahl S (2017). "Leveraging Uncertainty Information from Deep Neural Networks for Disease Detection". *Scientific Reports*. 7 (1): 17816. doi:[10.1038/s41598-017-17876-z](https://doi.org/10.1038/s41598-017-17876-z).
33. <sup>△</sup>Lee K, Lee H, Lee K, Shin J (2018). "Training Confidence-calibrated Classifiers for Detecting Out-of-Distribution Samples". arXiv:1711.09325. Available from: <https://arxiv.org/abs/1711.09325>.
34. <sup>△</sup><sup>a</sup>, <sup>△</sup><sup>b</sup>, <sup>△</sup><sup>c</sup> Mukhoti J, Gal Y (2018). "Evaluating bayesian deep learning methods for semantic segmentation". arXiv preprint arXiv:1811.12709. Available from: [arXiv:1811.12709](https://arxiv.org/abs/1811.12709).

35. <sup>a</sup>, <sup>b</sup>, <sup>c</sup> Mukhoti J, Kirsch A, van Amersfoort J, Torr PH, Gal Y (2023). "Deep Deterministic Uncertainty: A New Simple Baseline". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 24384--24394.
36. <sup>a</sup>, <sup>b</sup> Landgraf S, Hillemann M, Wursthorn K, Ulrich M (2024). "Uncertainty-aware Cross-Entropy for Semantic Segmentation". *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. 10: 129–136.
37. <sup>^</sup> Loquercio A, Segu M, Scaramuzza D (2020). "A general framework for uncertainty estimation in deep learning". *IEEE Robotics and Automation Letters*. 5 (2): 3153–3160.
38. <sup>a</sup>, <sup>b</sup>, <sup>c</sup> Wan S, Wu T-Y, Wong WH, Lee C-Y. Confnet: predict with confidence. 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE; 2018. p. 2921–2925.
39. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup> Nix DA, Weigend AS. Estimating the mean and variance of the target probability distribution. *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*. IEEE; 1994. 1:55–60.
40. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup>, <sup>e</sup> Gal Y, Ghahramani Z (2016). "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning." In: Balcan MF, Weinberger KQ, editors. *Proceedings of The 33rd International Conference on Machine Learning. Proceedings of Machine Learning Research*. 48:1050–1059. New York, New York, USA: PMLR. Available from: <https://proceedings.mlr.press/v48/gal16.html>. PDF: <http://proceedings.mlr.press/v48/gal16.pdf>.
41. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup>, <sup>e</sup> Valdenegro-Toro M (2023). "Sub-Ensembles for Fast Uncertainty Estimation in Neural Networks". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4119--4127.
42. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup>, <sup>e</sup> Lakshminarayanan B, Pritzel A, Blundell C (2017). "Simple and scalable predictive uncertainty estimation using deep ensembles." In: Guyon I, Luxburg U Von, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. *Advances in Neural Information Processing Systems*. 30. Curran Associates, Inc.
43. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup> Ayhan MS, Berens P (2018). "Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks". *Medical Imaging with Deep Learning*.
44. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup>, <sup>e</sup> Tyree S, Tremblay J, To T, Cheng J, Mosier T, Smith J, Birchfield S (2022). "6-DoF Pose Estimation of Household Objects for Robotic Manipulation: An Accessible Dataset and Benchmark". *International Conference on Intelligent Robots and Systems (IROS)*.
45. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>, <sup>d</sup>, <sup>e</sup> Nex F, Stathopoulou EK, Remondino F, Yang MY, Madhuanand L, Yogender Y, Alsadik B, Weinmann M, Jutzi B, Qin R (2024). "UseGeo-A UAV-based multi-sensor dataset for geospatial research". *ISPRS Open Journal of Photogrammetry and Remote Sensing*. 2024: 100070.

46. <sup>△</sup>Eigen D, Puhrsch C, Fergus R (2014). "Depth map prediction from a single image using a multi-scale deep network". *Advances in neural information processing systems*. 27.
47. <sup>△</sup>Ming Y, Meng X, Fan C, Yu H (2021). "Deep learning for monocular depth estimation: A review". *Neurocomputing*. 438: 14–33.
48. <sup>△</sup>Masoumian A, Rashwan HA, Cristiano J, Asif MS, Puig D (2022). "Monocular depth estimation using deep learning: A review". *Sensors*. 22 (14): 5353.
49. <sup>△</sup>Khan F, Salahuddin S, Javidnia H (2020). "Deep learning-based monocular depth estimation methods—a state-of-the-art review". *Sensors*. 20 (8): 2272.
50. <sup>△</sup>Arampatzakis V, Pavlidis G, Mitianoudis N, Papamarkos N (2023). "Monocular depth estimation: A thorough review". *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
51. <sup>△</sup>Fu H, Gong M, Wang C, Batmanghelich K, Tao D (2018). "Deep ordinal regression network for monocular depth estimation". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018–2011.
52. <sup>△</sup>Yuan W, Gu X, Dai Z, Zhu S, Tan P (2022). "Neural window fully-connected crfs for monocular depth estimation". *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 3916–3925.
53. <sup>△</sup>Patil V, Sakaridis C, Liniger A, Van Gool L. P3depth: Monocular depth estimation with a piecewise planarity prior. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022:1610–1621.
54. <sup>△</sup>Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*. 2020.
55. <sup>△</sup>Agarwal A, Arora C. Depthformer: Multiscale vision transformer for monocular depth estimation with global local information fusion. *2022 IEEE International Conference on Image Processing (ICIP)*. IEEE; 2022. p. 3873–3877.
56. <sup>△</sup>Li Z, Chen Z, Liu X, Jiang J (2023). "Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation". *Machine Intelligence Research*. 20 (6): 837–854.
57. <sup>△</sup>Yang G, Tang H, Ding M, Sebe N, Ricci E (2021). "Transformer-based attention networks for continuous pixel-wise prediction". *Proceedings of the IEEE/CVF International Conference on Computer vision*. 16269–16279.



58. <sup>^</sup>Aich S, Vianney JMU, Islam MA, Liu MKB. Bidirectional attention network for monocular depth estimation. 2021 IEEE International Conference on Robotics and Automation (ICRA). IEEE; 2021. p. 11746–11752.
59. <sup>^</sup>Zhao C, Zhang Y, Poggi M, Tosi F, Guo X, Zhu Z, Huang G, Tang Y, Mattoccia S. Monovit: Self-supervised monocular depth estimation with a vision transformer. 2022 international conference on 3D vision (3D V). IEEE; 2022. p. 668–678.
60. <sup>^</sup>Ning J, Li C, Zhang Z, Wang C, Geng Z, Dai Q, He K, Hu H (2023). "All in tokens: Unifying output space of visual tasks via soft token". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 19900--19910.
61. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>Ranftl R, Bochkovskiy A, Koltun V (2021). "Vision transformers for dense prediction". *Proceedings of the IEEE/CVF international conference on computer vision*. 12179--12188.
62. <sup>^</sup>Eftekhari A, Sax A, Malik J, Zamir A (2021). "Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10786--10796.
63. <sup>^</sup>Yin W, Zhang J, Wang O, Niklaus S, Mai L, Chen S, Shen C (2021). "Learning to recover 3d scene shape from a single image". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 204--213.
64. <sup>^</sup>Piccinelli L, Sakaridis C, Yu F (2023). "idisc: Internal discretization for monocular depth estimation". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21477--21487.
65. <sup>a</sup>, <sup>b</sup>Ke B, Obukhov A, Huang S, Metzger N, Dauterbach RC, Schindler K. Repurposing diffusion-based image generators for monocular depth estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024:9492-9502.
66. <sup>^</sup>Bhat SF, Alhashim I, Wonka P (2021). "Adabins: Depth estimation using adaptive bins". *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4009--4018.
67. <sup>^</sup>Li Z, Wang X, Liu X, Jiang J (2024). "Binsformer: Revisiting adaptive bins for monocular depth estimation". *IEEE Transactions on Image Processing*.
68. <sup>^</sup>Duan Y, Guo X, Zhu Z (2023). "Diffusiondepth: Diffusion denoising approach for monocular depth estimation". *arXiv preprint arXiv:2303.05021*.
69. <sup>^</sup>Ji Y, Chen Z, Xie E, Hong L, Liu X, Liu Z, Lu T, Li Z, Luo P. Ddp: Diffusion model for dense visual prediction. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023:21741-21752.
70. <sup>^</sup>Saxena S, Herrmann C, Hur J, Kar A, Norouzi M, Sun D, Fleet DJ (2024). "The surprising effectiveness of diffusion models for optical flow and monocular depth estimation". *Advances in Neural Information Proc*

essing Systems. 36.

71. <sup>^</sup>Saxena S, Kar A, Norouzi M, Fleet DJ (2023). "Monocular depth estimation using diffusion models". arXiv preprint arXiv:2302.14816.
72. <sup>^</sup>Patni S, Agarwal A, Arora C (2024). "ECoDepth: Effective Conditioning of Diffusion Models for Monocular Depth Estimation". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 28285--28295.
73. <sup>^</sup>Kumar CS, Bhandarkar SM, Prasad M (2018). "Monocular depth prediction using generative adversarial networks". *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 300--308.
74. <sup>^</sup>Aleotti F, Tosi F, Poggi M, Mattoccia S (2018). "Generative adversarial networks for unsupervised monocular depth prediction". *Proceedings of the European conference on computer vision (ECCV) workshops*. 0--0.
75. <sup>^</sup>Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, Bernstein MS, Bohg J, Bosselut A, Brunskill E, et al. *On the opportunities and risks of foundation models*. arXiv preprint arXiv:2108.07258. 2021.
76. <sup>^</sup>Guizilini V, Vasiljevic I, Chen D, Ambruş R, Gaidon A (2023). "Towards zero-shot scale-aware monocular depth estimation". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9233--9243.
77. <sup>^</sup>Yin W, Zhang C, Chen H, Cai Z, Yu G, Wang K, Chen X, Shen C (2023). "Metric3d: Towards zero-shot metric 3d prediction from a single image". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9043--9053.
78. <sup>a</sup> <sup>b</sup>MacKay DJC (1992). "A Practical Bayesian Framework for Backpropagation Networks". *Neural Computation*. 4 (3): 448-472. doi:10.1162/neco.1992.4.3.448.
79. <sup>a</sup> <sup>b</sup>Van Amersfoort J, Smith L, Teh YW, Gal Y (2020). "Uncertainty estimation using a single deep deterministic neural network". *International conference on machine learning*. PMLR, 9690--9700.
80. <sup>a</sup> <sup>b</sup>Liu J, Lin Z, Padhy S, Tran D, Bedrax Weiss T, Lakshminarayanan B (2020). "Simple and principled uncertainty estimation with deterministic deep learning via distance awareness". *Advances in Neural Information Processing Systems*. 33: 7498-7512.
81. <sup>^</sup>Amini A, Schwarting W, Soleimany A, Rus D (2020). "Deep evidential regression". *Advances in Neural Information Processing Systems*. 33: 14927-14937.
82. <sup>^</sup>Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R (2014). "Dropout: A Simple Way to Prevent Neural Networks from Overfitting". *Journal of Machine Learning Research*. 15 (56): 1929-1958.

83. <sup>^</sup>Ovadia Y, Fertig E, Ren J, Nado Z, Sculley D, Nowozin S, Dillon J, Lakshminarayanan B, Snoek J. Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. In: Wallach H, Larochelle H, Beygelzimer A, d'Alché-Buc F, Fox E, Garnett R, editors. *Advances in Neural Information Processing Systems*. 32nd ed. Curran Associates, Inc.; 2019. Available from: <https://proceedings.neurips.cc/paper/2019/file/8558cb408c1d76621371888657d2eb1d-Paper.pdf>.
84. <sup>^</sup>Wursthorn K, Hillemann M, Ulrich M (2022). "Comparison of Uncertainty Quantification Methods for CNN-based Regression". *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. XLIII-B2-2022: 721–728.
85. <sup>^</sup>Gustafsson FK, Danelljan M, Schon TB (2020). "Evaluating scalable bayesian deep learning methods for robust computer vision". *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. 318–319.
86. <sup>a</sup> <sup>b</sup> <sup>c</sup>Landgraf S, Hillemann M, Kapler T, Ulrich M (2024). "Evaluation of Multi-task Uncertainties in Joint Semantic Segmentation and Monocular Depth Estimation". *arXiv e-prints*. arXiv:--2405.
87. <sup>^</sup>Fort S, Hu H, Lakshminarayanan B. Deep Ensembles: A Loss Landscape Perspective. *arXiv*. 2020. Available from: [arxiv.org/abs/1912.02757](https://arxiv.org/abs/1912.02757).
88. <sup>a</sup> <sup>b</sup>Poggi M, Aleotti F, Tosi F, Mattoccia S. (2020). "On the uncertainty of self-supervised monocular depth estimation." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3227--3237.
89. <sup>^</sup>Hirose N, Taguchi S, Kawano K, Koide S (2021). "Variational monocular depth estimation for reliability prediction". *2021 International Conference on 3D Vision (3DV)*. IEEE, 637--647.
90. <sup>a</sup> <sup>b</sup>Nie X, Shi D, Li R, Liu Z, Chen X (2021). "Uncertainty-aware self-improving framework for depth estimation". *IEEE Robotics and Automation Letters*. 7(1): 41--48.
91. <sup>a</sup> <sup>b</sup>Choi H, Lee H, Kim S, Kim S, Kim S, Sohn K, Min D (2021). "Adaptive confidence thresholding for monocular depth estimation". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12808–12818.
92. <sup>^</sup>Dikov G, van Vugt J. Variational depth networks: Uncertainty-aware monocular self-supervised depth estimation. *European Conference on Computer Vision*. Springer; 2022. p. 43–60.
93. <sup>^</sup>Kendall A, Gal Y (2017). "What uncertainties do we need in bayesian deep learning for computer vision?" *Advances in neural information processing systems*. 30.
94. <sup>^</sup>Yu X, Franchi G, Aldea E (2021). "Slurp: Side learning uncertainty for regression problems". *arXiv preprint arXiv:2110.11182*.

95. <sup>△</sup>Hornauer J, Belagiannis V (2022). "Gradient-based uncertainty for monocular depth estimation". *European Conference on Computer Vision, Springer*, 613–630.
96. <sup>△</sup>Franchi G, Yu X, Bursuc A, Aldea E, Dubuisson S, Filliat D (2022). "Latent discriminant deterministic uncertainty". *European Conference on Computer Vision, Springer*, 243–260.
97. <sup>△</sup>Mi L, Wang H, Tian Y, He H, Shavit NN (2022). "Training-free uncertainty estimation for dense regression: Sensitivity as a surrogate". *Proceedings of the AAAI Conference on Artificial Intelligence*. 36 (9): 10042–10050.
98. <sup>△</sup>Oquab M, Darcet T, Moutakanni T, Vo H, Szafraniec M, Khalidov V, Fernandez P, Haziza D, Massa F, El-Nouby A, et al. *Dinov2: Learning robust visual features without supervision*. *arXiv preprint arXiv:2304.07193*. 2023.
99. <sup>△</sup>Yun S, Han D, Oh SJ, Chun S, Choe J, Yoo Y (2019). "Cutmix: Regularization strategy to train strong classifiers with localizable features". *Proceedings of the IEEE/CVF international conference on computer vision*. 6023--6032.
100. <sup>△</sup><sup>Ⓜ</sup>Wang S, Leroy V, Cabon Y, Chidlovskii B, Revaud J (2024). "Dust3r: Geometric 3d vision made easy". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20697--20709.
101. <sup>△</sup>Loshchilov I (2017). "Decoupled weight decay regularization". *arXiv preprint arXiv:1711.05101*.
102. <sup>△</sup>Hoda\{u\}n T, Sundermeyer M, Drost B, Labb\{u\}e Y, Brachmann E, Michel F, Rother C, Matas J (2020). "BOP challenge 2020 on 6D object localization". *Computer Vision--ECCV 2020 Workshops: Glasgow, UK, August 23--28, 2020, Proceedings, Part II 16*. Springer, 577--594.
103. <sup>△</sup>Sundermeyer M, Hoda\{u\}n T, Labbe Y, Wang G, Brachmann E, Drost B, Rother C, Matas J (2023). "Bop challenge 2022 on detection, segmentation and pose estimation of specific rigid objects". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2785--2794.
104. <sup>△</sup>Hodan T, Sundermeyer M, Labbe Y, Nguyen VN, Wang G, Brachmann E, Drost B, Lepetit V, Rother C, Matas J. *BOP Challenge 2023 on Detection Segmentation and Pose Estimation of Seen and Unseen Rigid Objects*. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024:5610–5619.
105. <sup>△</sup>Landgraf S, Hillemann M, Kapler T, Ulrich M. *Efficient multi-task uncertainties for joint semantic segmentation and monocular depth estimation*. *DAGM German Conference on Pattern Recognition (GCPR)*. Springer; 2024.
106. <sup>△</sup>Fort S, Hu H, Lakshminarayanan B (2019). "Deep ensembles: A loss landscape perspective". *arXiv preprint arXiv:1912.02757*.

## **Declarations**

**Funding:** The authors acknowledge support by the state of Baden-Württemberg through bwHPC. This work is supported by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT partition. The work is partially supported by the Office of Naval Research (N00014-23-1-2670).

**Potential competing interests:** No potential competing interests to declare.