

Single Image Depth Prediction Made Better: A Multivariate Gaussian Take

Ce Liu¹ Suryansh Kumar^{1*} Shuhang Gu² Radu Timofte^{1,3} Luc Van Gool^{1,4}
¹CVL ETH Zürich ²UESTC China ³University of Würzburg ⁴KU Leuven
 {ce.liu, sukumar, vangool}@vision.ee.ethz.ch
 shuanggu@uestc.edu.cn, radu.timofte@uni-wuerzburg.de

Abstract

Neural-network-based single image depth prediction (SIDP) is a challenging task where the goal is to predict the scene’s per-pixel depth at test time. Since the problem, by definition, is ill-posed, the fundamental goal is to come up with an approach that can reliably model the scene depth from a set of training examples. In the pursuit of perfect depth estimation, most existing state-of-the-art learning techniques predict a single scalar depth value per-pixel. Yet, it is well-known that the trained model has accuracy limits and can predict imprecise depth. Therefore, an SIDP approach must be mindful of the expected depth variations in the model’s prediction at test time. Accordingly, we introduce an approach that performs continuous modeling of per-pixel depth, where we can predict and reason about the per-pixel depth and its distribution. To this end, we model per-pixel scene depth using a multivariate Gaussian distribution. Moreover, contrary to the existing uncertainty modeling methods—in the same spirit, where per-pixel depth is assumed to be independent, we introduce per-pixel covariance modeling that encodes its depth dependency w.r.t. all the scene points. Unfortunately, per-pixel depth covariance modeling leads to a computationally expensive continuous loss function, which we solve efficiently using the learned low-rank approximation of the overall covariance matrix. Notably, when tested on benchmark datasets such as KITTI, NYU, and SUN-RGB-D, the SIDP model obtained by optimizing our loss function shows state-of-the-art results. Our method’s accuracy (named MG) is among the top on the KITTI depth-prediction benchmark leaderboard¹.

1. Introduction

Recovering the depth of a scene using images is critical to several applications in computer vision [2, 15, 25, 29, 30]. It is well founded that precise estimation of scene depth

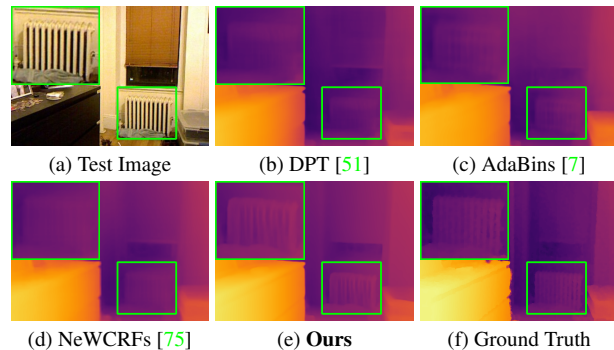


Figure 1. **Qualitative Comparison.** By modeling scene depth as multivariate Gaussian and enforcing the parametric low-rank covariance constraints in the loss function, we observe that our model can reliably predict depth for both high-frequency and low-frequency scene details. In the above example, we can notice better qualitative results than the state-of-the-art methods.

from images is likely only under multi-view settings [65]—which is indeed a correct statement and hard to contend². But what if we could effectively learn scene depth using images and their ground-truth depth values, and be able to predict the scene depth using just a single image at test time? With the current advancements in deep learning techniques, this seems quite possible empirically and has also led to excellent results for the single image depth prediction (SIDP) task [40, 51]. Despite critical geometric arguments against SIDP, practitioners still pursue this problem not only for a scientific thrill but mainly because there are several real-world applications in which SIDP can be extremely beneficial. For instance, in medical [42], augmented and virtual reality [21, 55], gaming [19], novel view synthesis [56, 57], robotics [64], and related vision applications [24, 51].

Regardless of remarkable progress in SIDP [1, 36, 37, 39, 40, 50, 75], the recent state-of-the-art deep-learning methods, for the time being, just predict a single depth value per pixel at test time [37]. Yet, as is known, trained models have accuracy limits. As a result, for broader adoption of SIDP in

*Corresponding Author (k.sur46@gmail.com)

¹http://www.cvlibs.net/datasets/kitti/eval_depth.php?benchmark=depth_prediction

²As many 3D scene configurations can have the same image projection.

applications, such as robot vision and control, it is essential to have information about the reliability of predicted depth. Consequently, we model the SIDP task using a continuous distribution function. Unfortunately, it is challenging, if not impossible, to precisely model the continuous 3D scene. In this regard, existing methods generally resort to increasing the size and quality of the dataset for better scene modeling and improve SIDP accuracy. On the contrary, little progress is made in finding novel mathematical modeling strategies and exploiting the prevalent scene priors. To this end, we propose a multivariate Gaussian distribution to model scene depth. In practice, our assumption of the Gaussian modeling of data is in consonance with real-world depth data (see Fig. 2) and generalizes well across different scenes. Furthermore, many computer and robot vision problems have successfully used it and benefited from Gaussian distribution modeling in the past [9, 20, 45, 53, 63, 72, 77].

Let’s clarify this out way upfront that this is not for the first time an approach with a motivation of continuous modeling for SIDP is proposed [4, 22, 26, 28, 31, 47]. Yet, existing methods in this direction model depth per pixel independently. It is clearly unreasonable, in SIDP modeling, to assume absolute democracy among each pixel, especially for very closeby scene points. Therefore, it is natural to think of modeling this problem in a way where depth at a particular pixel can help infer, refine, and constrain the distribution of depth value for other image pixels. Nevertheless, it has yet to be known a priori the neighboring relation among pixels in the scene space to define the depth covariance among them. We do that here by defining a very general covariance matrix of dimension `number of pixels × number of pixels`, *i.e.*, depth prediction at a given pixel is assumed to be dependent on all other pixels’ depth.

Overall, we aim to advocate multivariate Gaussian modeling with a notion of depth dependency among pixels as a useful scene prior. Now, add a fully dependent covariance modeling proposal to it—as suitable relations among pixels are not known. This makes the overall loss function computationally expensive. To efficiently optimize the proposed formulation, we parameterize the covariance matrix, assuming that it lies in a rather low-dimensional manifold so that it can be learned using a simple neural network. For training our deep network, we utilize the negative log likelihood as the loss function (cf. Sec. 3.1). The trained model when tested on standard benchmark datasets gives state-of-the-art results for SIDP task (see Fig. 1 for qualitative comparison).

Contributions. To summarize, our key contributions are:

- A novel formulation to perform multivariate Gaussian covariance modeling for solving the SIDP task in a deep neural network framework is introduced.
- The introduced multivariate Gaussian covariance modeling for SIDP is computationally expensive. To solve it efficiently, the paper proposes to learn the low-rank co-

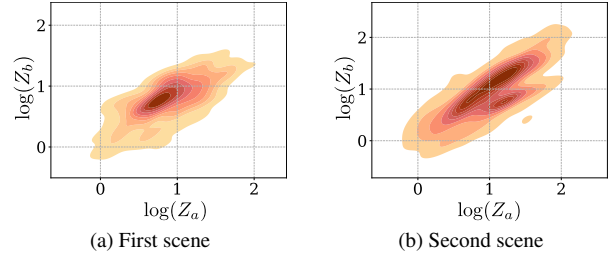


Figure 2. The marginal ground-truth depth distribution for a pixel pair Z_a, Z_b for two scenes. The depth values for the pixel pair are taken from the fixed image location in the dataset, but the selected images are visually similar for the suitability of the feature and its corresponding depth values. The statistics show that the Gaussian distribution assumption with covariance modeling is a sensible choice for SIDP problem and not an unorthodox belief arranged or staged for an intricate formulation.

- variance matrix approximation by deep neural networks.
- Contrary to the popular SIDP methods, the proposed approach provides better depth as well as a measure of the suitability of the predicted depth value at test time.

2. Related Work

Predicting the scene depth from a single image is a popular problem in computer vision with long-list of approaches. To keep our review of the existing literature succinct and on-point, we discuss work of direct relevance to the proposed method. Roughly, we divide well-known methods into two sub-category based on their depth prediction modeling.

(i) General SIDP. By general SIDP, we mean methods that predict one scalar depth value per image pixel at inference time. Earlier works include Markov Random Field (MRF) or Conditional Random Fields (CRF) modeling [58–60, 67]. With the advancement in neural network-based approaches, such classical modeling ideas are further improved using advanced deep-network design [35, 41, 70, 71, 75]. A few other stretches along this line use piece-wise planar scene assumption [33]. Other variations in deep neural network-based SIDP methods use ranking, ordinal relation constraint, or structured-guided sampling strategy [10, 14, 38, 69, 78]. The main drawback with the above deep-learning methods is that they provide an over-smoothed depth solution, and most of them rely on some heuristic formulation for depth-map refinement as a post-refinement step.

Recently, transformer networks have been used for better feature aggregation via an increase in the network receptive field [7, 8, 73, 75] or with the use of attention supervision [8] leading to better SIDP accuracy. Another mindful attempt is to exploit the surface normal and depth relation. To this end, [23] introduces both normal and depth loss for SIDP, whereas [74] proposes using virtual normal loss for imposing explicit 3D scene constraint and utilizing

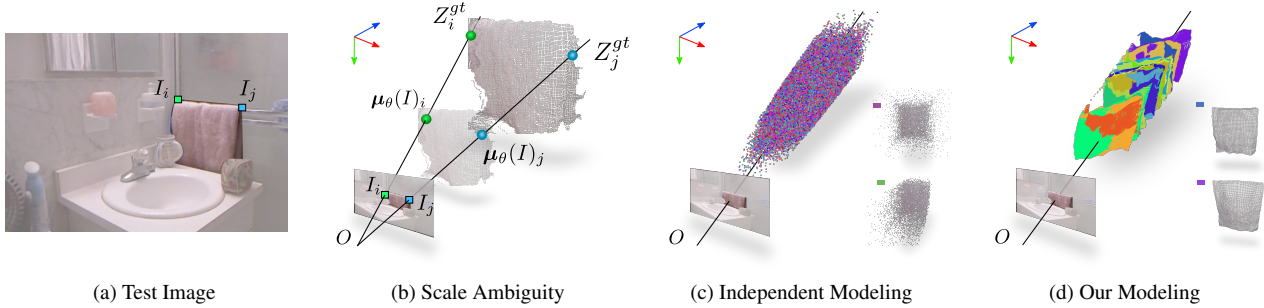


Figure 3. (a) Test image of an indoor bathroom scene. (b) The problem of scale ambiguity: showing several possible 3D point-cloud configurations of the towel with same imaging region. (c) If the depth values from the towel region are back-projected in the scene space under the independent Gaussian distribution assumption of the depth map. Clearly, the 3D point cloud results are not encouraging. (d) Samples due to our multivariate Gaussian distribution modeling that constrain the pixel depth with learned covariance. We observe the samples drawn from our modeling provide better 3D point clouds. Note: depth map is transformed to point cloud for visualization.

long-range scene dependencies. Long *et al.* [44] improved over [74] by introducing an adaptive strategy to compute local patch surface normal by randomly sampling for candidate triplets. A short time ago, [6] showed commendable results using normal-guided depth propagation [49] with a depth-to-normal and learnable normal-to-depth module.

(ii) Probabilistic SIDP. In comparison, there are limited pieces of literature that directly target to solve SIDP in a probabilistic way, where the methods could predict the scene depth and simultaneously can reason about its prediction quality. Generally, popular methods from uncertainty modeling in deep-network are used as it is for such purposes. For instance, Kendall *et al.* [26] Bayesian uncertainty in deep-networks, Lakshminarayanan *et al.* [31] deep ensemble-based uncertainty modeling, Amini *et al.* [4] deep evidential regression approach, is shown to work also for the depth prediction task. Yet, these methods are very general and can be used for most, if not all, computer vision problems [16]. Moreover, these methods treat each pixel independently, which may lead to inferior SIDP modeling.

This brings us to the point that application-specific priors, constraints, and settings could be exploited to enhance the solution, and we must not wholly depend on general frameworks to tackle the problem with similar motivation [4, 26, 31]. Therefore, this paper advocate using per-pixel multivariate Gaussian covariance modeling with efficient low-rank covariance parametric representation to improve SIDP for its broader application. Furthermore, we show that the depth likelihood due to multivariate Gaussian distribution modeling can help define better loss function and allow depth covariance learning based on scene feature regularity. With our modeling formulation, the derived loss function naturally unifies the essence of L2 loss, scale-invariant loss, and gradient loss. These three losses can be derived as a special case of our proposed loss (cf. Sec. 3.2).

3. Proposed Method

To begin with, let’s introduce problem setting and some general notations, which we shall be using in the rest of the paper. That will help concisely present our formulation, its useful mathematical insights, discussion, and application to Bayesian uncertainty estimation in deep networks.

Problem Setting. Given an image $\mathbf{I} \in \mathbb{R}^{m \times n}$ at test time, our goal is to predict the reliable per-pixel depth map $\mathbf{Z} \in \mathbb{R}^{m \times n}$, where m, n symbolize the number of image rows and cols, respectively. For this problem, we reshape the image and corresponding ground-truth depth map as a column vector represented as $I \in \mathbb{R}^{N \times 1}$ and $Z^{gt} \in \mathbb{R}^{N \times 1}$, respectively. Here, $N = m \times n$ is the total number of pixels in the image and \mathcal{D} denotes the train set.

3.1. Multivariate Gaussian Modeling

Let’s assume the depth map Z corresponding to image I follows a N -dimensional Gaussian distribution. Accordingly, we can write the distribution Φ given I as

$$\Phi(Z|\theta, I) = \mathcal{N}(\mu_\theta(I), \Sigma_\theta(I, I)). \quad (1)$$

Where, $\mu_\theta(I) \in \mathbb{R}^{N \times 1}$ and $\Sigma_\theta(I, I) \in \mathbb{R}^{N \times N}$ symbolize the mean and covariance of multivariate Gaussian distribution \mathcal{N} of predicted depth, respectively. The θ represents the parametric description of mean and covariance, which the neural network can learn at train time. It is important to note that with such network modeling, it is easy for the network to reliably reason about the scene depth distribution of similar-looking images at test time. Using the general form of multivariate Gaussian density function, the log probability density of Eq.(1) could be elaborated as

$$\log \Phi(Z|\theta, I) = -\frac{N}{2} \log 2\pi - \frac{1}{2} \log \det(\Sigma_\theta(I, I)) - \frac{1}{2} (Z - \mu_\theta)^T (\Sigma_\theta(I, I))^{-1} (Z - \mu_\theta). \quad (2)$$

Eq.(2) is precisely the formulation we strive to implement. Yet, computing the determinant and inverse of a $N \times N$ covariance matrix can be computationally expensive, *i.e.*, $O(N^3)$, for a reasonable image resolution. Previous methods in this direction usually restrict covariance to be diagonal [26, 31], *i.e.*, $\Sigma_\theta(I, I) = \mathbf{diag}(\sigma_\theta^2(I))$ with σ_θ as the standard deviation learned by the network with parameter θ . Even though such a simplification leads to computationally tractable algorithm $O(N)$, it leads to questionable depth prediction at test time. The reason for that is in SIDP, each pixel’s depth could vary by a single scale value which must be the same for all the pixels under the rigid scene assumption (see Fig.3b). By assuming covariance to be zero, each pixel is modeled independently; hence the coherence among scene points is lost completely. It can also be observed from Fig.(3c) when the covariance matrix is restricted to a diagonal matrix, the sampled depth from $\Phi(Z|\theta, I)$ is incoherently scattered. Therefore, it is pretty clear that multivariate covariance modeling is essential (see Fig. 3d) despite being computationally expensive.

To overcome the computational bottleneck in covariance modeling, we propose to exploit Σ_θ parametric form with low-rank assumption. It is widely studied in statistics that multivariate data relation generally has low-dimensional structure [32, 66, 76]. Since, covariance matrix is symmetric and positive definite, we write Σ_θ in parametric form *i.e.*,

$$\Sigma_\theta(I, I) = \Psi_\theta(I)\Psi_\theta(I)^T + \sigma^2\mathbf{eye}(N), \quad (3)$$

where, $\Psi_\theta(I) \in \mathbb{R}^{N \times M}$ is learned by deep networks with parameter θ with $M \lll N$. $\mathbf{eye}(N) \in \mathbb{R}^{N \times N}$ is a slang for identity matrix. $\Psi_\theta(I)\Psi_\theta(I)^T$ is symmetric and $\sigma^2\mathbf{eye}(N)$ guarantees positive definite matrix with $\sigma > 0$ as some positive constant. By using the popular matrix inversion lemma [48] and Eq.(3) parametric form, log probability density defined in Eq.(2) can be re-written as

$$\begin{aligned} \log \Phi(Z|\theta, I) = & -\frac{N}{2} \log 2\pi\sigma^2 - \frac{1}{2} \log \det(\mathbf{A}) - \\ & \frac{\sigma^{-2}}{2} \mathbf{r}^T \mathbf{r} + \frac{\sigma^{-4}}{2} \mathbf{r}^T \Psi_\theta(I) \mathbf{A}^{-1} \Psi_\theta(I)^T \mathbf{r}, \end{aligned} \quad (4)$$

with $\mathbf{r} = Z - \mu_\theta(I)$, and $\mathbf{A} = \sigma^{-2}\Psi_\theta(I)^T \Psi_\theta(I) + \mathbf{eye}(M)$. It can be shown that the above form for modeling covariance is computationally tractable with complexity $O(NM + M^3)$ [53] as compared to $O(N^3)$ since $M \lll N$ (refer to supplementary material for details). We use Eq.(4) as negative log likelihood (*NLL*) loss function, *i.e.*, $\mathcal{L}_{NLL} = -\log \Phi(Z^{gt}|\theta, I)$ to train the network for learning per-pixel depth, and covariance w.r.t. all the pixels, hence overcoming the shortcomings with prior works in SIDP.

3.2. Deeper Insights into the Formulation

A detailed analysis of Eq.(4) and how it naturally encapsulates the notion of popular loss functions are presented for

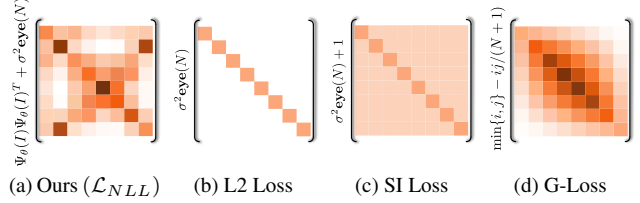


Figure 4. **The covariance matrix of loss function.** (a) Ours Σ_θ . (b)-(d) shows the equivalent covariance matrix for the (b) L2 loss, (c) scale invariant loss, and (d) gradient loss. It can be observed that our covariance already contains most, if not all, information that could be recovered by employing different loss functions, hence showing the generality of our formulation.

better understanding. Concretely, we show that “L2 Loss”, “Scale Invariant Loss (SI Loss)”, and “Gradient Loss (G-Loss)” as a special case of Eq.(4); thus, our formulation is more general. Later, in the subsection, we apply Eq.(4) to well-known Bayesian uncertainty modeling in deep neural networks showing improved uncertainty estimation than independent Gaussian assumption.

3.2.1 Relation to Popular Loss Function

By taking Eq.(4) *NLL* form as the training loss, *i.e.*, $-\log \Phi(Z^{gt}|\theta, I)$, we show that using the special values for $\Psi_\theta(I)$, the *NLL* loss can be reduced to some widely-used losses (see Fig. 4). Here, symbolizes $Z^{gt} \in \mathbb{R}^{N \times 1}$. Denoting $\mathbf{r} = Z^{gt} - \mu_\theta(I)$, we derive the relation.

(i) **Case I.** Substituting $\Psi_\theta(I) = \mathbf{0}_N$ in Eq.(4) will give

$$-\log \Phi(Z^{gt}|\theta, I) \propto \mathbf{r}^T \mathbf{r}, \quad (5)$$

which is equivalent to the “L2 loss” function. Here, $\mathbf{0}_N$ is a column vector with N elements, all set to 0.

(ii) **Case II.** Substituting $\Psi_\theta(I) = \mathbf{1}_N$ in Eq.(4) will give

$$-\log \Phi(Z^{gt}|\theta, I) \propto \mathbf{r}^T \mathbf{r} - \frac{\alpha}{N} (\mathbf{r}^T \mathbf{1}_N)^2, \quad (6)$$

where, $\alpha = (\sigma^{-2}N)/(\sigma^{-2}N + 1)$ and $\mathbf{1}_N$ is a column vector with N elements set to 1. Assuming $\sigma^{-2}N \ggg 1$, which is mostly the experimental setting in SIDP, then $\alpha \approx 1$ and Eq.(6) becomes equivalent to “Scale Invariant Loss”.

(iii) **Case III.** Here, we want to show the relation between Eq.(4) and gradient loss. But, unlike previous cases, it’s a bit involved. So, for simplicity, consider the gradient of the flattened depth map (*i.e.*, a column vector)³. The general squared gradient loss between the ground-truth and predicted depth can be computed as $(\nabla Z^{gt} - \nabla \mu_\theta(I))^T (\nabla Z^{gt} - \nabla \mu_\theta(I))$, where $\nabla \in \mathbb{R}^{N \times N}$ is the gradient operator for computing the first order difference of

³ignoring border pixels for simple 1D case

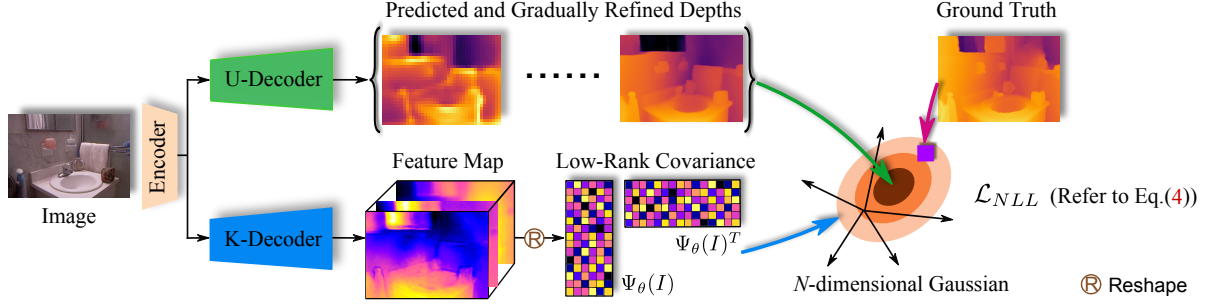


Figure 5. **Overview of our framework.** Given an image, first an encoder is employed to extract features. Then the U-Decoder will predict and gradually refine the depth maps. And the K-Decoder is responsible for predicting the factor $\Psi_\theta(I)$ for modeling the covariance. In the end, we compute the negative log likelihood of the N -dimensional Gaussian distribution as the loss function to supervise training.

Z^{gt} and $\mu_\theta(I)$ [54]. Taking out the common factor, we can re-write the gradient loss as $(\nabla(Z^{gt} - \mu_\theta(I)))^T (\nabla(Z^{gt} - \mu_\theta(I)))$. Simplifying using the matrix transpose property and it can be written in compact form as $\mathbf{r}^T (\nabla^T \nabla) \mathbf{r}$, which is equivalent to the Gaussian multivariate form in Eq.(2). Let's denote $J \triangleq (\nabla^T \nabla)^{-1}$, where $J_{i,j} = \min\{i, j\} - ij / (N + 1)$ [11]. However, J is difficult to parameterize and decompose into low-dimensional form. Concretely, we want to factorize J into $\Psi_\theta(I)\Psi_\theta(I)^T + \sigma^2 \mathbf{eye}(N)$ that fits the notion developed in Eq.(2) and Eq.(3), with $\Psi_\theta(I) \in \mathbb{R}^{N \times M}$ and $M \lll N$.

Fortunately, it is possible to approximate J as $J \approx (\Psi_\theta(I)\Psi_\theta(I)^T + \sigma^2 \mathbf{eye}(N))$ by using well-known Eigen approximation [18]. To be precise, setting $\Psi_\theta(I)$ to

$$\Psi_\theta(I)_{k,l} = \sqrt{\lambda(J)_l} \mathbf{U}(J)_{k,l} \quad (7)$$

where $\lambda(J) \in \mathbb{R}^{N \times 1}$ and $\mathbf{U}(J) \in \mathbb{R}^{N \times N}$ are the sorted eigenvalues and corresponding eigenvectors of J , respectively that can be computed using $\lambda(J)_l = (2 - 2 \cos \frac{l\pi}{N+1})^{-1}$ and $\mathbf{U}(J)_{k,l} = (-1)^{k+l} \sin \frac{kl\pi}{N+1}$ [46].

3.2.2 Application in Uncertainty Estimation

We apply Eq.(4) to the popular Bayesian uncertainty modeling in neural networks. Given $\Phi(Z|\theta, I)$ as aleatoric uncertainty for depth map Z [26], we can compute the Bayesian uncertainty by marginalising over the parameters' posterior distribution using the following well-known equation:

$$\Phi(Z|I, \mathcal{D}) = \int \Phi(Z|\theta, I) \Phi(\theta|\mathcal{D}) d\theta \quad (8)$$

where \mathcal{D} is the train set. The analytic integration of Eq.(8) is difficult to compute in practice, and is usually approximated by Monte Carlo integration, such as ensemble [31] and dropout [16]. Suppose we have sampled a set of parameters $\Theta \triangleq \{\theta^s\}_{s=1}^S$ from $\Phi(\theta|\mathcal{D})$. The integration is popularly approximated as

$$\Phi(Z|I, \mathcal{D}) = \frac{1}{S} \sum_s \Phi(Z|\theta^s, I). \quad (9)$$

The $\Phi(Z|I, \mathcal{D})$ denotes the mixture of Gaussian distributions [31]. The mean and covariance of the distribution is computed as $\bar{\mu}(I) = \frac{1}{S} \sum_s \mu^s(I)$ and $\bar{\Sigma}(I, I) = \bar{\Psi}(I)\bar{\Psi}(I)^T + \sigma^2 \mathbf{eye}(N)$, respectively [68], which in fact has the same form as Eq.(3), where we compute $\bar{\Psi}$ using the following expression

$$\bar{\Psi} = \frac{1}{\sqrt{S}} \text{concat}(\Psi^1, \dots, \Psi^S, \mu^1 - \bar{\mu}, \dots, \mu^S - \bar{\mu}). \quad (10)$$

So, from the derivations in Sec.(3.2.1) and Sec.(3.2.2), it is quite clear that our proposed Eq.(4) is more general and encapsulates flavors of popular loss functions widely used in deep networks. For the SIDP problem we need such a loss function for deep neural network parameters learning. Next, we discuss the implementation in our proposed pipeline and usefulness of our introduced loss function.

3.3. Overall Pipeline

To keep our pipeline description simple, let's consider the image and depth map in 2D form instead of a column vector. For brevity, we slightly abuse the notation hereafter. Here, we use the same notation we defined for the 1D Gaussian distribution case for simplicity. For a better understanding of our overall pipeline, we provide architectural implementation details following Fig.(5) blueprint, *i.e.*, (i) Encoder details, (ii) Decoder details, followed by (iii) Train and test time settings.

(i) Encoder Details. Our encoder takes the image $\mathbf{I} \in \mathbb{R}^{m \times n}$ as input and gives a hierarchical feature maps $\mathbf{F} = \{\mathbf{F}^1, \mathbf{F}^2, \mathbf{F}^3, \mathbf{F}^4\}$ as output, where $\mathbf{F}^i \in \mathbb{R}^{C^i \times m^i \times n^i}$ denotes a feature map with channels C^i , and resolution $m^i \times n^i$. We adopt the Swin-Large [43] as the encoder. Specifically, it includes four stages of non-linear transformation to extract features from the input image, where each stage contains a series of transformer blocks to learn non-linear transformation and a downsampling block to reduce the resolution of the feature map by 2. We collect the output feature map from the last block of the i -th stage as \mathbf{F}_i .

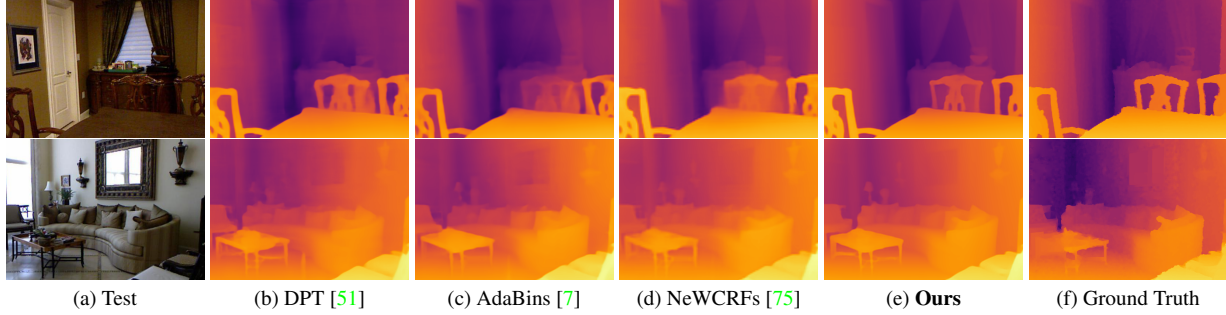


Figure 6. **Qualitative Comparison.** Our method recovers better depth even for complex scenes. Our depth results even qualitatively looks closer to the ground truth than the prior art such as (b) DPT [51], (c) AdaBins [7], (d) NeWCRFs [75] on NYU Depth V2 test set [61].

(ii) **Decoder Details.** The *U-decoder* (see Fig.5) estimates a set of depth maps $\{\mu_{\theta}^i(\mathbf{I})\}_{i=1}^4$. The U-decoder first estimates $\mu_{\theta}^4(\mathbf{I})$ only from \mathbf{F}^4 by a convolution layer, then upsamples and refines the depth map in a hierarchical manner. At the i -th stage, where i decreases from 3 to 1, we first concatenate $\mu_{\theta}^{i+1}(\mathbf{I})$ and \mathbf{F}^{i+1} from the previous stage, and then feed into a stack of convolution layers to refine the depth map and feature map. The refined depth map is upsampled via bi-linear interpolation to double the resolution, and denoted as $\mu_{\theta}^i(\mathbf{I})$. Similarly, the refined feature map is upsampled and added to \mathbf{F}^i . In the end, we upsample all the depth maps $\{\mu_{\theta}^i(\mathbf{I})\}_{i=1}^4$ into $m \times n$ resolution via bi-linear interpolation as the final output of the U-decoder.

The *K-decoder* (see Fig. 5) estimates $\Psi_{\theta}(\mathbf{I})$. It first upsamples and refines the feature maps in \mathbf{F} . Specifically, at the i -th stage, where i decreases from 3 to 1, it upsamples \mathbf{F}^{i+1} from the previous stage and adds to the \mathbf{F}^i . We utilize a stack of convolution layers to further refine the added feature map. In the end, we upsample the refined feature map \mathbf{F}^1 to $m \times n$ resolution by bi-linear interpolation, and predict the $\Psi_{\theta}(\mathbf{I})$ by a convolution layer.

(iii) **Train and Test Time Setting.** At train time, we collect $\{\mu_{\theta}^i(\mathbf{I})\}_{i=1}^4$, and $\Psi_{\theta}(\mathbf{I})$ and compute loss using our proposed loss function (refer to Sec. 3.4). At test time, our approach provides $\mu_{\theta}^1(\mathbf{I})$ as the final depth map prediction. Furthermore, we can query $\Psi_{\theta}(\mathbf{I})$ to infer the distribution of the depth map if necessary depending on the application.

3.4. Loss Function

As shown in Sec. 3.2.1, the negative log likelihood loss can approximate the scale invariant loss and the gradient loss when $\Psi_{\theta}(I)$ and σ take special values. Consequently, we propose the following overall loss function:

$$\mathcal{L}_{total} = \sum_{j=1}^4 \mathcal{L}_{NLL}^j + \frac{1}{N} \sum_i (\mu_{\theta}^1(I)_i - Z_i^{gt})^2 \quad (11)$$

where \mathcal{L}_{NLL}^j is the negative log likelihood loss applying to $\mu_{\theta}^j(I)$ and $\Psi_{\theta}(I)$. Note, however, Eq.(11) second term is optional. Yet, it is added to provide train time improvement.

4. Experiments and Results

Implementation Details. We implemented our framework in PyTorch 1.7.1 and Python 3.8 with CUDA 11.0. All the experiments and statistical results shown in the draft are simulated on a computing machine with Quadro RTX 6000 (24GB Memory) GPU support. We use evaluation metrics including SILog, Abs Rel, RMS, RMS log, δ_i , Sq Rel, iRMS to report our results on the benchmark dataset. For exact definition of the metrics we refer to [33].

Datasets. We performed experiments and statistical comparisons with the prior art on benchmark datasets such as NYU Depth V2 [61], KITTI [17], and SUN RGB-D [62].

(a) **NYU Depth V2:** This dataset contains images of indoor scenes with 480×640 resolution [61]. We follow the standard train and test split setting used by previous works for experiments [33]. Precisely, we use 24,231 image-depth pairs for training the network and 654 images for testing the performance. Note that the depth map evaluation for this dataset has an upper bound of 10 meters.

(b) **KITTI:** This dataset contains images and depth data of outdoor driving scenarios. The official experimental split contains 42,949 training images, 1,000 validation images, and 500 test images with 352×1216 resolution [17]. Here, the depth map accuracy can be evaluated up to an upper bound of 80 meters. In addition, there are few works following the split from Eigen [13], which includes 23,488 images for training and 697 images for the test.

(c) **SUN RGB-D:** It contains data of indoor scenes captured by different cameras [62]. The depth values range from 0 up to 10 meters. The images are resized to 480×640 resolution for consistency. We use the official test set [62] of 5050 images to evaluate the generalization of the frameworks.

Training Details. We use Adam optimizer [27] to minimize our proposed loss function and learn network parameters. At train time, the learning rate is decreased from $3e^{-5}$ to $1e^{-5}$ by the cosine annealing scheduler. Our encoder—which is inspired from [43], is initialized by pre-training the network on ImageNet [12]. For the KITTI dataset, we

Method	Backbone	SILog ↓	Abs Rel ↓	RMS ↓	RMS log ↓	δ_1 ↑
GeoNet [49]	ResNet-50	-	0.128	0.569	-	0.834
DORN [14]	ResNet-101	-	0.115	0.509	-	0.828
VNL [74]	ResNeXt-101	-	0.108	0.416	-	0.875
TransDepth [73]	ViT-B	-	0.106	0.365	-	0.900
ASN [44]	HRNet-48	-	0.101	0.377	-	0.890
BTS [33]	DenseNet-161	11.533	0.110	0.392	0.142	0.885
DPT-Hybrid [51]	ViT-B	9.521	0.110	0.357	0.129	0.904
AdaBins [7]	EffNet-B5+ViT-mini	10.570	0.103	0.364	0.131	0.903
ASTrans [8]	ViT-B	10.429	0.103	0.374	0.132	0.902
NeWCRFs [75]	Swin-L	9.102	0.095	0.331	0.119	0.922
Ours	Swin-L	8.323	0.087	0.311	0.110	0.933
% Improvement		8.56%	8.42%	6.04%	7.56%	1.18%

Table 1. Comparison with the state-of-the-art methods on the NYU test set [61]. Please refer to Sec. 4.1 for details.

Method	SILog ↓	Abs Rel ↓	Sq Rel ↓	iRMS ↓
DLE [39]	11.81	9.09	2.22	12.49
DORN [14]	11.80	8.93	2.19	13.22
BTS [33]	11.67	9.04	2.21	12.23
BANet [3]	11.55	9.34	2.31	12.17
PWA [34]	11.45	9.05	2.30	12.32
ViP-DeepLab [50]	10.80	8.94	2.19	11.77
NeWCRFs [75]	10.39	8.37	1.83	11.03
Ours	9.93	7.99	1.68	10.63
% Improvement	4.43%	4.54%	8.20%	3.63%

Table 2. Comparison with the state-of-the-art methods on the KITTI official test set [17]. We only list the results from the published methods.

train our framework for 10 and 20 epochs on the official split [17] and Eigen [13] split, respectively. For the NYU dataset [61], our framework is trained for 20 epochs. We randomly apply horizontal flipping on the image and depth map pair at train time for data augmentation.

4.1. Performance Comparison with Prior Works

Tab. 1, Tab. 2, Tab. 3, and Tab. 4 show our method’s statistical performance comparison with popular state-of-the-art (SOTA) methods on NYU Depth V2 [61], KITTI official [17] and Eigen [13] split, and SUN RGB-D [62]. From the tables, it is easy to infer that our approach consistently performs better on all the popular evaluation metrics. The percentage improvement over the previous SOTA is indicated in green for better exposition. In particular, on the NYU test set, which is a challenging dataset, we reduce the SILog error from the previous best result of 9.102 to 8.323 and increase δ_1 metric from 0.922 to 0.933⁴. Fig. 6 shows qualitative comparison results. It can be observed that our method’s predicted depth is better at low and high-frequency scene details. For the SUN RGB-D test set, all competing models, including ours, are trained on the NYU DepthV2 train set without fine-tuning on SUN RGB-D [62]. In addition, we align the predictions from all the models with the ground truth by a scale and shift following [52]. Tab. 4 results indicate our method’s better generalization capability than other approaches. More results are provided in the supplementary material.

⁴At the time of submission, our method’s performance on the KITTI official leaderboard was the best among all the published works.

4.2. Bayesian Uncertainty Estimation Comparison

In this part, we compare with the classical Bayesian dropout [26], which uses independent Gaussian distribution to quantify uncertainty. As for our approach, we also use dropout to sample multiple depth predictions, and compute the negative log likelihood following the distribution in Eq.(9). More specifically, in each block of Swin transformer [43], we randomly drop feature channels before the layer normalization [5] operation with probability 0.01. We first sample $S = 10$ predictions for each test image, then compute the mean and covariance of the mixture of Gaussian distributions in Eq.(9), and further approximate the entire distribution as single Gaussian following [31]. We present the comparison results of the negative log likelihood in Fig. 7. Our multivariate Gaussian distribution achieves much lower negative log likelihood cost.

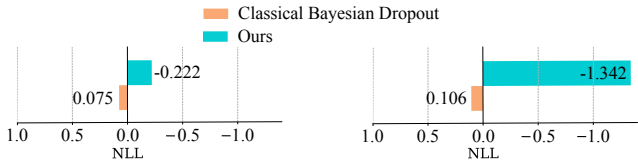


Figure 7. Comparison with the classical Bayesian dropout for uncertainty estimation. The left and right figures present the Negative Log Likelihood (NLL) of the predicted depth map distribution on KITTI Eigen [13,17] split and NYU test set [61] respectively. Our multivariate Gaussian distribution achieves lower NLL than the independent Gaussian distribution in classical Bayesian dropout.

4.3. Ablations and Further Analysis

To better understand our introduced approach, we performed ablations on the NYU Depth V2 dataset [61] and studied our trained model’s inference time and memory footprint for its practical suitability.

(i) **Effect of NLL Loss.** To realize the significance of NLL loss in Eq.(11), we replaced it with L_2 loss, SI loss [13], gradient loss, and virtual normal loss [74] one by one, while keeping the remaining term in Eq.(11) fixed. The statistical results are shown in Tab. 5. The stats show that our proposed NLL loss achieves the best performance over all the widely used metrics.

Method	Backbone	SILog ↓	Abs Rel ↓	RMS ↓	RMS log ↓	δ_1 ↑
DORN [14]	ResNet-101	-	0.072	0.273	0.120	0.932
VNL [74]	ResNeXt-101	-	0.072	0.326	0.117	0.938
TransDepth [73]	ViT-B	8.930	0.064	0.275	0.098	0.956
BTS [33]	DenseNet-161	8.933	0.060	0.280	0.096	0.955
DPT-Hybrid [51]	ViT-B	8.282	0.062	0.257	0.092	0.959
AdaBins [7]	EffNet-B5+ViT-mini	8.022	0.058	0.236	0.089	0.964
ASTrans [8]	ViT-B	7.897	0.058	0.269	0.089	0.963
NeWCRFs [75]	Swin-L	6.986	0.052	0.213	0.079	0.974
Ours	Swin-L	6.757	0.050	0.202	0.075	0.976
% Improvement		3.28%	3.85%	5.16%	5.06%	0.21%

Table 3. Comparison with the state-of-the-art methods on the KITTI Eigen test set [13]. Please refer to Sec. 4.1 for details.

Method	Backbone	SILog ↓	Abs Rel ↓	RMS ↓	RMS log ↓	δ_1 ↑
AdaBins [7]	EffNet-B5+ViT-mini	13.652	0.110	0.321	0.137	0.906
NeWCRFs [75]	Swin-L	13.695	0.105	0.322	0.138	0.920
Ours	Swin-L	11.985	0.090	0.282	0.120	0.936
% Improvement		12.49%	14.29%	12.42%	13.04%	1.74%

Table 4. Comparison with AdaBins [7] and NeWCRFs [75] on SUN RGB-D test set [62]. All methods are trained on NYU Depth V2 train set without fine-tuning on SUN RGB-D. Please refer to Sec. 4.1 for details.

Loss	SILog ↓	Abs Rel ↓	RMS ↓	δ_1 ↑
L2	8.912	0.090	0.324	0.929
SI [13]	8.762	0.089	0.322	0.929
Gradient	8.886	0.090	0.323	0.929
VNL [74]	8.543	0.090	0.325	0.926
Ours	8.323	0.087	0.311	0.933

Table 5. Comparison of our NLL loss function with widely used loss functions for solving SIDP task.

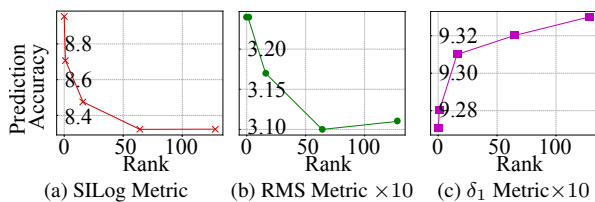


Figure 8. Depth prediction accuracy of our method using different evaluation metrics w.r.t change in the rank of the covariance matrix. The increase in the rank improves prediction accuracy and shows saturation at 128, thereby showing the effectiveness of our low-dimensional modeling.

(ii) **Performance with the change in the Rank of Covariance.** We vary the rank of $\Psi_\theta(I)$, and observe our method’s prediction accuracy. We present the accuracy under various evaluation metrics in Fig. 8. With increase in the rank, the distribution is better approximated, and the performance improves, but saturates later showing the suitability of its low-dimensional representation.

(iii) **Evaluation on NeWCRFs.** We evaluate our loss function on NeWCRFs [75] network design using their proposed training strategies. The depth prediction accuracy is shown in Tab. 6. The results convincingly indicate the benefit of our proposed loss on a different SIDP network.

Method	SILog ↓	Abs Rel ↓	RMS ↓	δ_1 ↑
NeWCRFs	9.102	0.095	0.331	0.922
+Our Loss	8.619	0.086	0.316	0.935

Table 6. Results using our loss on [75] network on NYU Depth.

(iv) **Inference Time & Parameter Comparison.** We compared our method’s inference time, and the number of model parameters to the recent state-of-the-art NeWCRFs [75]. The inference time is measured on the NYU Depth V2 test set [61] with batch size 1. As shown in Tab. 7, our method achieves lower scale invariant logarithmic error (SILog) with fewer network parameters and comparable FPS. Such a statistical performance further endorse our approach’s practical adequacy.

Method	SILog ↓	Speed (FPS) ↑	Param (M) ↓
NeWCRFs	9.171	10.551	258
Ours	8.323	9.909	244

Table 7. Comparison of the inference time and parameters with NeWCRFs [75] on NYU Depth V2 [61].

Refer supplementary for more experiments and results

5. Conclusion

This work suitably formalizes the connection between robust statistical modeling techniques, *i.e.*, multivariate covariance modeling with low-rank approximation, and popular loss functions in neural network-based SIDP problem. The novelty presented in this paper arises from the fact that the proposed pipeline and loss term turns out to be more general, hence could be helpful in the broader application of SIDP in several tasks, such as depth uncertainty for robot vision, control and others. Remarkably, the proposed formulation is not only theoretically compelling but observed to be practically beneficial, resulting in a loss function that is used to train the proposed network showing state-of-the-art SIDP results on several benchmark datasets.

Acknowledgments. This work was partly supported by ETH General Fund (OK), Chinese Scholarship Council (CSC), and The Alexander von Humboldt Foundation.

References

- [1] Ashutosh Agarwal and Chetan Arora. Attention attention everywhere: Monocular depth prediction with skip attention. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5861–5870, 2023.
- [2] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011.
- [3] Shubhra Aich, Jean Marie Uwabeza Vianney, Md Amirul Islam, and Mannat Kaur Bingbing Liu. Bidirectional attention network for monocular depth estimation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11746–11752. IEEE, 2021.
- [4] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *Advances in Neural Information Processing Systems*, 33:14927–14937, 2020.
- [5] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [6] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Iron-depth: Iterative refinement of single-view depth using surface normal and its uncertainty. In *British Machine Vision Conference (BMVC)*, 2022.
- [7] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021.
- [8] Wenjie Chang, Yueyi Zhang, and Zhiwei Xiong. Transformer-based monocular depth estimation with attention supervision. In *32nd British Machine Vision Conference (BMVC 2021)*, 2021.
- [9] Duo Chen, Zixin Tang, Zhenyu Xu, Yunan Zheng, and Yiguang Liu. Gaussian fusion: Accurate 3d reconstruction via geometry-guided displacement interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5916–5925, 2021.
- [10] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. *Advances in neural information processing systems*, 29, 2016.
- [11] C.M. da Fonseca and J. Petronilho. Explicit inverses of some tridiagonal matrices. *Linear Algebra and its Applications*, 325(1):7–21, 2001.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [13] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014.
- [14] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2002–2011, 2018.
- [15] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [16] Yarin Gal et al. Uncertainty in deep learning.
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [18] Gene H Golub and Charles F Van Loan. *Matrix computations*. JHU press, 2013.
- [19] Mohammad Mahdi Haji-Esmaili and Gholamali Montazer. Playing for depth. *arXiv preprint arXiv:1810.06268*, 2018.
- [20] Lukas Hewing, Juraj Kabzan, and Melanie N Zeilinger. Cautious model predictive control using gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6):2736–2743, 2019.
- [21] Derek Hoiem, Alexei A Efros, and Martial Hebert. Automatic photo pop-up. In *ACM SIGGRAPH 2005 Papers*, pages 577–584. Association for Computing Machinery, 2005.
- [22] Julia Hornauer and Vasileios Belagiannis. Gradient-based uncertainty for monocular depth estimation. In *European Conference on Computer Vision*, pages 613–630. Springer, 2022.
- [23] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1043–1051. IEEE, 2019.
- [24] Nishant Jain, Suryansh Kumar, and Luc Van Gool. Robustifying the multi-scale representation of neural radiance fields. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022.
- [25] Berk Kaya, Suryansh Kumar, Carlos Oliveira, Vittorio Ferrari, and Luc Van Gool. Uncertainty-aware deep multi-view photometric stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12601–12611, 2022.
- [26] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017.
- [27] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [28] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: learning sfm from sfm. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 698–713, 2018.
- [29] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames. In *Proceedings of the IEEE international conference on computer vision*, pages 4649–4657, 2017.
- [30] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Superpixel soup: Monocular dense 3d reconstruction of a complex dynamic scene. *IEEE transactions on pattern analysis and machine intelligence*, 43(5):1705–1717, 2019.

- [31] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- [32] Miguel Lázaro-Gredilla, Joaquin Quinonero-Candela, Carl Edward Rasmussen, and Aníbal R Figueiras-Vidal. Sparse spectrum gaussian process regression. *The Journal of Machine Learning Research*, 11:1865–1881, 2010.
- [33] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019.
- [34] Sihaeng Lee, Janghyeon Lee, Byungju Kim, Eojindl Yi, and Junmo Kim. Patch-wise attention network for monocular depth estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1873–1881, 2021.
- [35] Bo Li, Chunhua Shen, Yuchao Dai, Anton van den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [36] Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation. *arXiv preprint arXiv:2203.14211*, 2022.
- [37] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *arXiv preprint arXiv:2204.00987*, 2022.
- [38] Julian Lienen, Eyke Hüllermeier, Ralph Ewerth, and Nils Nommensen. Monocular depth estimation via listwise ranking using the plackett-luce model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14595–14604, June 2021.
- [39] Ce Liu, Shuhang Gu, Luc Van Gool, and Radu Timofte. Deep line encoding for monocular 3d object detection and depth prediction. In *32nd British Machine Vision Conference (BMVC 2021)*, page 354, 2021.
- [40] Ce Liu, Suryansh Kumar, Shuhang Gu, Radu Timofte, and Luc Van Gool. Va-depthnet: A variational approach to single image depth prediction. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [41] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2024–2039, 2015.
- [42] Xingtong Liu, Ayushi Sinha, Masaru Ishii, Gregory D Hager, Austin Reiter, Russell H Taylor, and Mathias Unberath. Dense depth estimation in monocular endoscopy with self-supervised learning methods. *IEEE transactions on medical imaging*, 39(5):1438–1447, 2019.
- [43] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021.
- [44] Xiaoxiao Long, Cheng Lin, Lingjie Liu, Wei Li, Christian Theobalt, Ruigang Yang, and Wenping Wang. Adaptive surface normal constraint for depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12849–12858, 2021.
- [45] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [46] Silvia Noschese, Lionello Pasquini, and Lothar Reichel. Tridiagonal toeplitz matrices: properties and novel applications. *Numerical linear algebra with applications*, 20(2):302–326, 2013.
- [47] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. On the uncertainty of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3227–3237, 2020.
- [48] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, USA, second edition, 1992.
- [49] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. Geonet: Geometric neural network for joint depth and surface normal estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 283–291, 2018.
- [50] Siyuan Qiao, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Vip-deeplab: Learning visual perception with depth-aware video panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3997–4008, 2021.
- [51] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12179–12188, 2021.
- [52] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [53] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.
- [54] M. Renardy and R.C. Rogers. *An Introduction to Partial Differential Equations*. Texts in Applied Mathematics. Springer New York, 2004.
- [55] Stephan R Richter, Hassan Abu Al Haija, and Vladlen Koltun. Enhancing photorealism enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [56] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12216–12225, 2021.
- [57] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12892–12901, 2022.
- [58] Ashutosh Saxena, Sung Chung, and Andrew Ng. Learning depth from single monocular images. *Advances in neural information processing systems*, 18, 2005.

- [59] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Learning 3d scene structure from a single still image. In *2007 IEEE 11th international conference on computer vision*, pages 1–8. IEEE, 2007.
- [60] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE transactions on pattern analysis and machine intelligence*, 31(5):824–840, 2008.
- [61] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *European conference on computer vision*, pages 746–760. Springer, 2012.
- [62] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- [63] Ashley W Stroupe, Martin C Martin, and Tucker Balch. Merging gaussian distributions for object localization in multi-robot systems. In *Experimental Robotics VII*, pages 343–352. Springer, 2001.
- [64] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. Cnn-slam: Real-time dense monocular slam with learned depth prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6243–6252, 2017.
- [65] Shimon Ullman. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153):405–426, 1979.
- [66] Jiayi Wang, Raymond KW Wong, and Xiaoke Zhang. Low-rank covariance function estimation for multidimensional functional data. *Journal of the American Statistical Association*, 117(538):809–822, 2022.
- [67] Xiaoyan Wang, Chunping Hou, Liangzhou Pu, and Yonghong Hou. A depth estimating method from a single image using foe crf. *Multimedia Tools and Applications*, 74(21):9491–9506, 2015.
- [68] N.A. Weiss, P.T. Holmes, and M. Hardy. *A Course in Probability*. Pearson Addison Wesley, 2006.
- [69] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 611–620, 2020.
- [70] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [71] Dan Xu, Wei Wang, Hao Tang, Hong Liu, Nicu Sebe, and Elisa Ricci. Structured attention guided convolutional neural fields for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3917–3925, 2018.
- [72] Kohei Yamashita, Shohei Nobuhara, and Ko Nishino. 3d-gmnet: Single-view 3d shape recovery as A gaussian mixture. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020.
- [73] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformer-based attention networks for continuous pixel-wise prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16269–16279, 2021.
- [74] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5684–5693, 2019.
- [75] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Neural window fully-connected crfs for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3916–3925, June 2022.
- [76] Rui Zhou, Jiayi Ying, and Daniel P Palomar. Covariance matrix estimation under low-rank factor model with nonnegative correlations. *IEEE Transactions on Signal Processing*, 70:4020–4030, 2022.
- [77] Xiaosu Zhu, Jingkuan Song, Lianli Gao, Feng Zheng, and Heng Tao Shen. Unified multivariate gaussian mixture for efficient neural image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17612–17621, 2022.
- [78] Daniel Zoran, Phillip Isola, Dilip Krishnan, and William T Freeman. Learning ordinal relationships for mid-level vision. In *Proceedings of the IEEE international conference on computer vision*, pages 388–396, 2015.