

ON MONOCULAR DEPTH ESTIMATION AND UNCERTAINTY QUANTIFICATION USING CLASSIFICATION APPROACHES FOR REGRESSION

Xuanlong Yu^{1,2} Gianni Franchi² Emanuel Aldea¹

¹SATIE, Paris-Saclay University ²U2IS, ENSTA Paris, Institut Polytechnique de Paris

ABSTRACT

Monocular depth is important in many tasks, such as 3D reconstruction and autonomous driving. Deep learning based models achieve state-of-the-art performance in this field. A set of novel approaches for estimating monocular depth consists of transforming the regression task into a classification one. However, there is a lack of detailed descriptions and comparisons for Classification Approaches for Regression (CAR) in the community and no in-depth exploration of their potential for uncertainty estimation. To this end, this paper will introduce a taxonomy and summary of CAR approaches, a new uncertainty estimation solution for CAR, and a set of experiments on depth accuracy and uncertainty quantification for CAR-based models on KITTI dataset. The experiments reflect the differences in the portability of various CAR methods on two backbones. Meanwhile, the newly proposed method for uncertainty estimation can outperform the ensembling method with only one forward propagation.

Index Terms— Depth estimation, Uncertainty Estimation

1. INTRODUCTION

In machine learning, regression tasks predict a continuous output based on a given input. Yet, if the ground truth (prediction target) is within a specific range, *e.g.*, in the case of age estimation [1], one can quantize the ground truth and cast regression into a classification problem. We refer to these techniques as Classification Approaches for Regression (CAR), and in this paper we explore CAR techniques applied on monocular depth estimation.

Monocular depth estimation (MDE), which is an ill-posed problem [2], consists in predicting the scene depth given only an RGB image as the input. Deep Neural Networks (DNNs) learn the mapping between the single RGB images and their corresponding depth maps to solve MDE, and show good performance on indoor and outdoor benchmarks [3, 4].

Classification Approaches for Regression (CAR) [5, 6, 7, 8, 9, 10] have emerged recently in the spotlight among MDE algorithms. The core idea is to transfer regression to a classification problem using quantization (or discretization) strategies. The classification models can natively provide the con-

CAR-MDEs		Discretization			Adaptive	Post processing
		Handcrafted + One-hot	Handcrafted + Ordinal	Handcrafted + Smooth		
Loss function	KL divergence/	-	-	SORN [8]	-	Argmax
	Weighted CE loss	-	-	Cao et al. [5] ([15])	-	-
	CE loss	Li et al. [6] ([16, 17])	-	-	-	Soft weighted sum
	Multiple BCE loss	-	-	Yang et al. [9]	-	-
Post processing	Regression loss	DS-SIDENet [18]	-	-	Adabins [10]	-
	Ordinal	-	DORN [7] ([19, 20])	-	-	Ordinal sum
	Regression loss	-	-	-	-	-

Table 1: Summary on CAR-MDE solutions. In parentheses are methods that use the corresponding schemes as part of their solutions. The categories under Discretization, Loss function and Post processing are detailed in Sec.2.

fidence for prediction results, which also has the potential to improve the prediction accuracy [8]. DNNs are prone to two kinds of uncertainty: aleatoric uncertainty and epistemic uncertainty [11]. It is crucial to study the uncertainty of DNNs if we want to rely on their predictions. Some works proposed to estimate the uncertainty of MDE DNNs by using an auxiliary network [12, 13], or ensembling [14]. Here we gain access to the uncertainty directly using the CAR DNN.

This work will investigate and show the complete picture of the CAR MDE methods. The contributions are as follows:

1. We systematically summarize and formalize the all major CAR MDE mechanisms to the best of our knowledge;
2. We implement these mechanisms on top of two different backbones, comparing depth prediction and uncertainty quality on various evaluation metrics;
3. We propose a new, effective uncertainty estimation method named Expectation of Distance for CAR MDE models.

2. OVERVIEW OF CLASSIFICATION APPROACHES FOR REGRESSION (CAR)

2.1. Taxonomy of CAR MDE

To better unify the terms and make it easier to grasp the differences in contributions of CAR strategies, we propose to decompose the CAR problems into three key components: *discretization*, *loss function* and *post-processing*. Table 1 offers an overview of the specific strategies used in the previous works. The details are provided in the following sections.

The contributions of previous CAR based MDE solutions fall in two main groups: 1. novel strategies in the three components mentioned above [5, 6, 7, 8, 9, 10]; 2. architecture and/or loss modifications based on the previous strategies [16, 19, 17, 20, 15, 18]. In most papers, CAR can improve model accuracy, making it outperform its regression version [5, 6, 7, 8, 9, 10, 16], or may improve model per-

We acknowledge the support of the Saclay-IA computing platform.

formance as part of multi-task learning [18, 17].

2.1.1. General notations

Let us first consider a monocular depth dataset $D = \{(\mathbf{x}_i, \mathbf{d}_i)\}_i$, where $\mathbf{x}_i \in \mathbb{R}^{3 \times H \times W}$, and $\mathbf{d}_i \in (\mathbb{R}^+)^{H \times W}$ represents the ground truth depth \mathbf{d}_i for the image \mathbf{x}_i . We denote $\{d_{i,j}\}_j^N$ all the pixel values in \mathbf{d}_i , where N is the number of pixels with valid ground truth. a, b are two real values representing the minimum and maximum depth value for the dataset.

For CAR strategies, we denote K the number of classes, which represents the level of discretization. Additionally, to simplify the notations, we use \log as the logarithm with base e for all papers except for [16, 18], where the \log refers to the logarithm with base 10. We denote f_{θ_1} the DNN with parameters θ_1 . Given \mathbf{x}_i , the prediction of f_{θ_1} :

$$\hat{\mathbf{y}}_i = \frac{e^{f_{\theta_1}(\mathbf{x}_i)}}{\sum_{p=0}^{c-1} e^{[f_{\theta_1}(\mathbf{x}_i)]_p}} \quad (1)$$

where $f_{\theta_1}(\mathbf{x}_i) \in \mathbb{R}^{c \times H \times W}$ is the logit map and $[f_{\theta_1}(\mathbf{x}_i)]_p$ its p -th coefficient, and $\hat{\mathbf{y}}_i$ is the Softmax output. The number of its channels $c = K$ by default, otherwise equals to the specific settings as in DORN [7] ($2K$) and Adabins [10] (128).

2.1.2. Discretization: Fully Handcrafted

The discretization function will output two components given \mathbf{d} : a *depth table* $\bar{\mathbf{d}} = \{\bar{d}_p\}_p^K \in \mathbb{R}^K$ where the possible discrete depth values are set ordinally, and an indicator map equivalent to a *classification map* $\mathbf{y}_i = \{\{y_{i,j,p}\}_p^K\}_j^N \in (\mathbb{R}^+)^{K \times H \times W}$ which points for each pixel the closest discrete depth value. This closest depth value can be considered as a class, leading to a classification task. Both $\bar{\mathbf{d}}$ and \mathbf{y}_i are *handcrafted*, and the goal becomes the learning of \mathbf{y}_i .

Handcrafted $\bar{\mathbf{d}}$: $\bar{\mathbf{d}}$ contains K values representing the centers of intervals $[\bar{d}_0, \bar{d}_1, \dots, \bar{d}_{K-1}, \bar{d}_K]$ with an interval width q :

$$\begin{aligned} \bar{\mathbf{d}} &= \{\bar{d}_p\}_p^K = \{(\bar{d}_0 + \bar{d}_1)/2, \dots, (\bar{d}_{K-1} + \bar{d}_K)/2\} \\ \bar{d}_k &= \log a + k \cdot q, k \in [0, K], q = (\log b - \log a)/K \end{aligned} \quad (2)$$

Handcrafted $\bar{\mathbf{d}}$ + One-hot \mathbf{y}_i : Given $\bar{\mathbf{d}}$, constructing \mathbf{y}_i is done using one-hot encoding, as applied in [6, 16, 17, 18]:

$$\mathbf{y}_{i,j}^{[\text{onehot}]} = [y_{i,j,0} \ \dots \ y_{i,j,k} \ \dots \ y_{i,j,K-1}] \in \mathbb{R}^K \quad (3)$$

with $y_{i,j,k} = 1$, if $k = \lfloor (\log(d_{i,j}/a))/q \rfloor$, and 0 otherwise.

in which, $\lfloor \cdot \rfloor$ is a rounding operator, q is defined in Eq. 2.

Handcrafted $\bar{\mathbf{d}}$ + Ordinal \mathbf{y}_i : Furthermore, there are several variants of Eq 3. Ordinal properties can be applied on it as presented in [7] and the followed works [19, 20]:

$$\mathbf{y}_{i,j}^{[\text{ord}]} = [y_{i,j,0} \ \dots \ y_{i,j,k} \ \dots \ y_{i,j,K-1}] \in \mathbb{R}^K \quad (4)$$

with $y_{i,j,k} = 1$, if $k \leq \lfloor (\log(d_{i,j}/a))/q \rfloor$, and 0 otherwise.

Handcrafted $\bar{\mathbf{d}}$ + Smooth \mathbf{y}_i : It is also possible to predict a smooth discrete map from the initial discrete map $\mathbf{y}_{i,j}^{[\text{onehot}]}$. The indicator in the classification map is softened by applying

on $\bar{\mathbf{d}}$ a Gaussian kernel, as to predict distance within a coarser range. The smooth \mathbf{y}_i are defined by:

$$\mathbf{y}_{i,j}^{[\text{smo1}]} = e^{-\gamma \|\log(d_{i,j}) - \bar{\mathbf{d}}\|^2} \quad (5)$$

$$\mathbf{y}_{i,j}^{[\text{smo2}]} = \frac{e^{-\gamma \|\log(d_{i,j}) - \bar{\mathbf{d}}\|^2}}{\sum_{p=0}^{K-1} e^{-\gamma \|\log(d_{i,j}) - \bar{d}_p\|^2}} \quad (6)$$

where γ is a hyperparameter which can be regarded as the scale of the discrete distribution (the smaller γ , the flatter the label distribution in $\mathbf{y}_{i,j}$). Specifically, Yang et al. [9] use Eq 5 as the unnormalized soft target labels, while SORN [8] applies the normalized version in Eq 6. Moreover, Cao et al. [5] introduce a $K \times K$ symmetric ‘‘information gain’’ matrix H in their loss function with elements $H(k, \mathbf{p}) = e^{-\gamma \|k - \mathbf{p}\|^2}$, where $\mathbf{p} = [0, \dots, K-1]$ and k is the discrete ground truth index as defined in Eq 3. In this case:

$$\mathbf{y}_{i,j}^{[\text{smo3}]} = e^{-\gamma \|k - \mathbf{p}\|^2} = e^{-\gamma \cdot q^{-2} \|\log(d_{i,j}) - (\bar{\mathbf{d}} - 0.5q)\|^2} \quad (7)$$

Since q is a constant, this strategy can be regarded as being equivalent to the $\mathbf{y}_{i,j}^{[\text{smo1}]}$ in Eq 5.

2.1.3. Discretization: Adaptive

In the absence of the handcrafted depth table or the classification map, one may also implicitly train both of them using a regression loss as in Adabins [10]. Thus the goal of the DNN is changed from fitting the handcrafted classification maps to fitting the continuous ground truth depth, while still following the principle of building depth tables and classification maps.

In Adabins [10], the depth table is implicitly trained along with the classification map using a non-linear block g_{θ_2} with θ_2 the parameters of g , which is a mini ViT [10]. In this case, g_{θ_2} is set on top of the backbone f_{θ_1} , and it will output $\hat{\mathbf{d}}_i^{[\text{ada}]}$ and $\hat{\mathbf{y}}_i^{[\text{ada}]}$ given $f_{\theta_1}(\mathbf{x}_i)$:

$$\hat{\mathbf{d}}_i^{[\text{ada}]} = \{a + (b - a) \left(\sum_{s=0}^p \hat{d}_{i,s}^{[\text{ada}]} \right)_{p=0}^{K-1}\}, \hat{\mathbf{y}}_i^{[\text{ada}]} = \frac{e^{l_{i,j}}}{\sum_{p=0}^{K-1} e^{l_{i,j,p}}} \quad (8)$$

$$\text{with } \{\hat{d}_{i,p}^{[\text{ada}]}\}_p^K, \mathbf{1}_i = g_{\theta_2}(f_{\theta_1}(\mathbf{x}_i))$$

where, $\hat{\mathbf{y}}_i^{[\text{ada}]}$ is the product of a Softmax function, and $\hat{\mathbf{d}}_i^{[\text{ada}]}$ is a cumulative summation output followed by a normalization operation which is included in g_{θ_2} . Since $\hat{\mathbf{d}}_i^{[\text{ada}]}$ is a product of g_{θ_2} taking $f_{\theta_1}(\mathbf{x}_i)$, for each \mathbf{x}_i , not only a unique classification map but also a unique depth table will be provided.

2.1.4. Loss function

Based on the previous discretization strategies, we introduce here the loss function design. Models should fit their output to the designed \mathbf{y} or \mathbf{d} . For brevity, we define first the total loss $L_{\text{total}} = \sum_i \sum_{j=0}^{N-1} L_{i,j}$ where $L_{i,j}$ is the loss for pixel j , on the i -th data. We just define $L_{i,j}$ in the following sections for simplicity.

Cross entropy (CE) loss: is a straightforward solution given an one-hot classification map: $L_{i,j}^{[\text{CE}]}(\theta_1) = -(\mathbf{y}_{i,j}^{[\text{onehot}]} \log \hat{\mathbf{y}}_{i,j})$.

Ordinal regression loss: is essentially an implicit ordinal selection plus a multiple binary cross entropy (BCE) loss. Instead of directly using $\hat{y}_{i,j}$, it requires to do an ordinal selection on the logit map $f_{\theta_1}(\mathbf{x}_i)$ with $c = 2K$ to $c = K$ as the predicted classification map, then to apply a Multiple-BCE loss on it:

$$L_{i,j}^{[\text{ord}]}(\theta_1) = -[\mathbf{y}_{i,j}^{[\text{ord}]} \log \hat{\mathbf{y}}_{i,j}^{[\text{ord}]} + (1 - \mathbf{y}_{i,j}^{[\text{ord}]}) \log(1 - \hat{\mathbf{y}}_{i,j}^{[\text{ord}]})] \quad (9)$$

with $\hat{\mathbf{y}}_{i,j}^{[\text{ord}]} = \frac{e^{[f_{\theta_1}(\mathbf{x}_i)]_{2p+1}}}{e^{[f_{\theta_1}(\mathbf{x}_i)]_{2p+1}} + e^{[f_{\theta_1}(\mathbf{x}_i)]_{2p}}}$

where $2p + 1$ and $2p$ represent the indices of the coefficient.

Weighted CE loss: is applied when the target vector is a soft discrete distribution. The CE loss turns to be equal to: $L_{i,j}^{[\text{WCE}]}(\theta_1) = -(\mathbf{y}_{i,j}^{[\text{smo2}]} \log \hat{\mathbf{y}}_{i,j})$, and it has the same form as the Kullback-Leibler divergence loss.

Multiple BCE loss: is another solution when the target is a soft discrete distribution. Yang et al. [9] apply BCE loss on every class value in $\mathbf{y}_{i,j}^{[\text{smo1}]}$ defined in Eq. 5.

Regression losses: DS-SIDENet [18] applies CAR with a smooth L1 loss [21] to fit the one-hot classification map target. Conversely, Adabins [10] combines the per-image adaptive depth table $\hat{\mathbf{d}}_i^{[\text{ada}]}$ and $\hat{\mathbf{y}}_{i,j}^{[\text{ada}]}$ defined in Eq. 8 as the predicted depth, then applies a Scale-Invariant loss [22].

2.1.5. Post-processing

Post-processing aims to restore the discrete predicted labels to continuous depth values. In the following equations, we use the power function in base e , see Sec. 2.1.1.

Ordinal sum: For DORN [7], the continuous depth is restored from the sum of the output Sigmoid labels which are higher than or equal to 0.5:

$$\hat{d}_{i,j} = \exp\{\log(a) + q \cdot [\sum_{p=0}^{K-1} \mathbb{1}\{\hat{y}_{i,j,p}^{[\text{ord}]} \geq 0.5\} + 0.5]\} \quad (10)$$

where q and $\hat{y}_{i,j,p}^{[\text{ord}]}$ are defined in Eq. 2 and Eq. 9 respectively.

Soft weighted sum: is a solution that may applied on both handcrafted or learned depth tables. It sums the Hadamard product between the depth table and the classification map:

$$\hat{d}_{i,j} = \exp\{\sum_{p=0}^{K-1} \bar{d}_p \cdot \hat{y}_{i,j,p}\} \quad (11)$$

Note that essentially AdaBins [10] also follows this pattern.

Argmax: The authors of SORN [8] claim that Argmax outperforms Soft weighted sum in their case:

$$\hat{d}_{i,j} = \exp\{\log(a) + q \cdot [\text{argmax}\{\{\hat{y}_{i,j,p}\}_p^K\} + 0.5]\} \quad (12)$$

where q is defined in Eq. 2.

2.2. Uncertainty estimation of CAR MDE

In this section, we will discuss the previous works on uncertainty estimation for CAR MDE, the difficulty of this problem and our proposed approaches on estimating CAR uncertainty.

The ground truth uncertainty or the Oracle should be the model's prediction error. The previous works on MDE uncertainty estimation [13, 12] mainly use the principle of learning the prediction error [11]. Meanwhile, the **Variance** among the point estimations given by MC-Dropout [23] and Deep Ensembles [14] can also be applied for this task. Unlike the previous works, the likelihoods of the predicted class (the quantified depth value) given the input data provided by CARs can offer another possibility to estimate the uncertainty mentioned in the previous works but rarely discussed. Yang et al. [9] suggest to use **Shannon Entropy (S-Entr)** [24] among the output Softmax classification map: $-\sum_{p=0}^{K-1} \hat{y}_{i,j,p} \log \hat{y}_{i,j,p}$. Moreover, they showed cases where the depth is well predicted, yet the entropy is high, leading to an under-confident uncertainty score. Other strategies such as **1-Maximum Class Probability (1-MCP)** can also be regarded as the uncertainty: $1 - \max_p(\{\hat{y}_{i,j,p}\}_p)$. These are typical solutions used in classification tasks, and we argue that they will be suitable in case of using Argmax in post-processing for CAR problems. Widely used soft weighted sum (see Table 1) makes the property of CAR special: not only the classification map but also the depth table should be taken into account in the final result as shown in Eq. 11.

Following these remarks, we propose a new solution for CAR uncertainty. We first note that the previously mentioned methods lack consideration of the depth table, and further its relationship to the classification map. Hence, we define as CAR uncertainty metric: the **Expectation of Distance (E-Dist)** between the quantified depth values (either handcrafted logarithm depth table $\bar{\mathbf{d}}$ (Eq. 2) or the learned one $\hat{\mathbf{d}}_i^{[\text{ada}]}$ (Eq. 8)) and the final predicted depth $\hat{\mathbf{d}}_i$ (Eq. 11, 12) :

$$\text{E-Dist} = \sum_{p=0}^{K-1} \hat{y}_{i,j,p} \cdot (e^{\bar{d}_p} - \hat{d}_{i,j})^2 \quad \text{or} \quad \text{E-Dist} = \sum_{p=0}^{K-1} \hat{y}_{i,j,p} \cdot (\hat{d}_{i,p}^{[\text{ada}]} - \hat{d}_{i,j})^2 \quad (13)$$

Additionally, to our knowledge, no previous works discuss the uncertainty of ordinal regression model [7]. According to its CAR strategy, only values greater than or equal to 0.5 in its classification map will be considered in the final depth calculation, thus we argue that the uncertainty comes from this part. The modified E-Dist for ordinal regression is as below: We propose to discretize the depth prediction $\hat{\mathbf{d}}_i$ (Eq. 10) using Eq. 4, that we denote as $\mathbf{y}_{i,j}^{[\text{ord}]}$. Then we calculate the distance between $\mathbf{y}_{i,j}^{[\text{ord}]}$ and $\hat{\mathbf{y}}_{i,j}^{[\text{ord}]}$ (defined in Eq.9) weighted by the depth table and only consider the part with $\hat{y}_{i,j}^{[\text{ord}]} \geq 0.5$:

$$\sum_{p=0}^{K-1} e^{\bar{d}_p} \cdot (y_{i,j,p}^{[\text{ord}]} - \hat{y}_{i,j,p}^{[\text{ord}]})^2 \cdot \mathbb{1}\{\hat{y}_{i,j,p}^{[\text{ord}]} \geq 0.5\} \quad (14)$$

3. EXPERIMENTS

In this section, we fill in the missing comparisons of the previous works. Meanwhile, our experiments provide an extensive analysis of CAR MDE uncertainty estimation. While it is not trivial to propose a model-agnostic approach, the ensuing discussion establishes some important guidelines about performing this task on CAR models.

Backbones	BTS							FCN							K		
	:11	:21	:31	Abs Rel	Sq RMSE	RMSE log ₁₀	log ₁₀	:11	:21	:31	Abs Rel	Sq RMSE	RMSE log ₁₀	log ₁₀			
DORN [7]	0.952	0.992	0.998	0.069	0.267	2.302	0.103	0.029	0.940	0.990	0.998	0.076	0.292	2.962	0.113	0.033	80
Cao et al. [5]	0.945	0.992	0.998	0.077	0.292	2.988	0.111	0.034	0.928	0.989	0.998	0.084	0.344	3.223	0.122	0.036	50
Li et al. [6]	0.950	0.990	0.998	0.070	0.287	2.928	0.106	0.030	0.940	0.988	0.997	0.075	0.314	3.190	0.116	0.033	150 [16]
SORN [8]	0.947	0.992	0.998	0.071	0.290	2.929	0.107	0.031	0.963	0.976	0.995	0.119	0.563	3.938	0.163	0.051	120
Yang et al. [9]	0.951	0.991	0.998	0.066	0.276	2.897	0.103	0.029	0.940	0.989	0.997	0.072	0.302	3.096	0.113	0.032	128
DS-SIDE [18]	0.950	0.991	0.998	0.071	0.275	2.886	0.106	0.032	0.931	0.990	0.998	0.079	0.331	3.353	0.119	0.035	80
Adabins [10]	0.935	0.990	0.996	0.078	0.347	3.143	0.114	0.033	0.937	0.991	0.998	0.079	0.331	3.027	0.113	0.033	256
DORN [7]	0.952	0.992	0.998	0.069	0.267	2.302	0.103	0.029	0.940	0.990	0.998	0.076	0.292	2.962	0.113	0.033	80
Cao et al. [5]	0.953	0.991	0.998	0.066	0.268	2.857	0.103	0.029	0.934	0.989	0.997	0.076	0.319	3.124	0.117	0.033	80
Li et al. [6]	0.949	0.990	0.997	0.087	0.305	2.982	0.116	0.037	0.933	0.988	0.997	0.096	0.350	3.157	0.126	0.040	80
SORN [8]	0.949	0.993	0.998	0.072	0.283	2.902	0.106	0.031	0.963	0.976	0.995	0.122	0.373	3.950	0.165	0.052	80
Yang et al. [9]	0.948	0.991	0.998	0.070	0.284	2.973	0.107	0.031	0.940	0.990	0.997	0.076	0.308	3.067	0.115	0.034	80
DS-SIDE [18]	0.950	0.991	0.998	0.071	0.275	2.886	0.106	0.032	0.931	0.990	0.998	0.079	0.331	3.353	0.119	0.035	80
Adabins [10]	0.933	0.989	0.998	0.079	0.357	3.203	0.116	0.033	0.937	0.990	0.998	0.076	0.318	3.062	0.112	0.032	80
Org	0.955	0.993	0.998	0.060	0.249	2.798	0.096	0.027	0.944	0.992	0.998	0.069	0.275	2.938	0.107	0.030	1
MC-Dropout [23]	0.941	0.992	0.998	0.083	0.308	2.910	0.114	0.035	0.918	0.984	0.996	0.085	0.369	3.157	0.125	0.036	1
Deep Ensembles [14]	0.957	0.993	0.999	0.059	0.233	2.688	0.093	0.026	0.946	0.992	0.998	0.068	0.269	2.923	0.106	0.030	1

(a)

Backbones	BTS						FCN						K
	AUSE RMSE			AUSE AbsRel			AUSE RMSE			AUSE AbsRel			
Methods	I-MCP	S-Entr	E-Dist	I-MCP	S-Entr	E-Dist	I-MCP	S-Entr	E-Dist	I-MCP	S-Entr	E-Dist	
Cao et al. [5]	0.542	0.770	0.133	0.382	0.424	0.411	0.532	0.701	0.127	0.354	0.375	0.375	50
Li et al. [6]	0.174	0.153	0.187	0.276	0.262	0.409	0.137	0.138	0.132	0.241	0.235	0.259	150
SORN [8]	1.371	1.394	0.157	0.939	0.982	0.427	1.244	1.283	0.170	0.754	0.755	0.451	120
Yang et al. [9]	0.141	0.145	0.094	0.232	0.219	0.256	0.142	0.161	0.111	0.225	0.226	0.247	128
DS-SIDE [18]	0.698	0.806	0.293	0.525	0.544	0.397	1.212	1.331	0.995	0.830	0.722	0.484	80
Adabins [10]	0.855	0.827	0.179	0.536	0.527	0.377	0.984	0.945	0.191	0.608	0.589	0.398	256
DORN [7]	0.188	0.158	0.128	0.530	0.430	0.303	0.202	0.165	0.135	0.593	0.445	0.283	80
Cao et al. [5]	0.371	0.476	0.119	0.323	0.329	0.356	0.349	0.393	0.117	0.284	0.265	0.306	80
Li et al. [6]	0.206	0.178	0.181	0.355	0.348	0.449	0.170	0.163	0.124	0.318	0.310	0.333	80
SORN [8]	1.367	1.390	0.157	0.900	0.941	0.444	1.228	1.275	0.175	0.725	0.737	0.473	80
Yang et al. [9]	0.194	0.179	0.099	0.273	0.259	0.271	0.156	0.169	0.104	0.258	0.252	0.274	80
DS-SIDE [18]	0.698	0.806	0.293	0.525	0.544	0.397	1.212	1.331	0.995	0.830	0.722	0.484	80
Adabins [10]	0.823	0.683	0.181	0.499	0.450	0.360	0.775	0.710	0.234	0.502	0.478	0.391	80
DORN [7]	0.188	0.158	0.128	0.530	0.430	0.303	0.202	0.165	0.135	0.593	0.445	0.283	80
MC-Dropout [23]	0.460	(Variance)	0.501	(Variance)	0.322	(Variance)	0.456	(Variance)	0.456	(Variance)	0.456	(Variance)	1
Deep Ensembles [14]	0.165	(Variance)	0.281	(Variance)	0.184	(Variance)	0.290	(Variance)	0.290	(Variance)	0.290	(Variance)	1

(b)

Table 2: The best/second-best values are highlighted in dark/light blue. The results from regression-based models are provided as reference (lower rows), and we only highlight the CAR-based results. (a) Depth uncertainty evaluations. MC-Dropout and Deep Ensembles will only provide the uncertainty with one method. (b) Depth accuracy evaluations. Org: the original BTS [25] model and the regression version applied on FCN [26] model.




Fig. 1: Experiment pipeline. Three heads will be applied on the backbone: **I:** original regression version; **II:** MDE with handcrafted discretization; **III:** MDE with adaptive discretization through a mini ViT module [10].

3.1. Experiment settings

All the experiments are based on Eigen-split [22] KITTI dataset [3]. We followed the original settings in the corresponding papers for CAR strategies and applied them on a regression-based and a classification-based backbone respectively. In addition, we added experiments with $K = 80$ to the methods with originally different choices for K for better comparison. Fig 1 illustrates the experiment pipeline.

Reg.-based backbone: We use BTS-DenseNet161 [25, 27]. Same BTS training settings are applied for all the methods.

Cla.-based backbone: We choose FCN-ResNet101 [26, 28]. FCN is originally designed for semantic segmentation, thus it is suitable for CAR methods. For the one-channel regression version (org), followed BTS, we apply a Sigmoid on the top and multiply the output by b .

Evaluation matrices: We use the same matrices first introduced in [22] and used in many subsequent works for depth performance. For uncertainty estimation, we use the area under sparsification error curve (AUSE), as in [9, 12, 13]. 1% of pixels are removed each time and we calculate RMSE and AbsRel for the rest. Uncertainty estimation methods we used are introduced in Sec. 2.2. We will compare the CAR MDE uncertainty with widely used MC-Dropout [23] (with 8 forward passes) and Deep Ensembles [14] (with 3 models).

Training time consumption: We use one NVIDIA Titan RTX to count the average time consumption on Forward+Backward passes for one image for all CAR methods with $K = 80$ as well as the original regression method and the Deep Ensembles [14] using the same training settings.

3.2. Performance and discussions

Table 2a and Table 2b provide depth and uncertainty results.

Depth: We find that all CAR MDE methods are portable, but training directly with the settings of the original backbones degrades performance. We discover that the Adabins, DS-SIDE and SORN [10, 18, 8] based models are more sensitive

CAR Solutions	Time consumption (ms)								
	DORN [7]	Cao et al. [5]	Li et al. [6]	SORN [8]	Yang et al. [9]	DS-SIDE [18]	Adabins [10]	Org	Deep Ensembles [14]
BTS	610.96	509.18	431.32	444.66	613.98	430.14	421.06	378.98	1136.94
FCN	735.38	614.52	538.92	556.84	722.80	540.70	588.26	517.66	1552.98

Table 3: Time consumption on Forward+Backward passes for one image.

to the selected backbone than the other ones. Despite the influence of the backbones, we also consider that the training settings for the original Adabins are more different from the ones of BTS. This difference may cause Adabins to produce worse performance after porting. DORN-based model [7] achieves the best result among CAR DNNs, which confirms the effectiveness of ordinal constraints.

Uncertainty: Our proposed E-Dist shows good and robust performance in most cases given a CAR MDE method. Among the CAR strategies, we found that the uncertainty quality is related to the sharpness of the labeling during discretization, and also to the loss function. Li et al. [6], Yang et al. [9] and Cao et al. [5] based DNNs perform better for the uncertainty. Li et al. [6] model has one-hot encoded labels in the classification map which leads to the sharpest label distribution. Yang et al. [9] model has $\gamma = 15$ in Eq. 5 and we can also have $\gamma \cdot q^{-2} = 65$ in Eq. 7 for Cao et al. [5] model. This big coefficient can sharpen the label distribution. Conversely, in SORN [8] the γ in Eq. 6 is much smaller, which results in the evenner distributed labels, and we consider this is the main cause of its worse performance. Yang et al. [9] based model outperforms the others, which indicates that the Multi-BCE loss is more suitable for uncertainty estimation, which is similar to the one-versus-all strategy [29].

Choices of K : According to two sets of results separated by K , the performance rankings are consistent: the depth precision and the uncertainty quality increase along K , and K has biggest impact on Cao et al. [5] based model.

Time efficiency: According to Table. 3, we argue that the CAR strategy will slightly slow down the training, especially for the ones requiring label smoothing in discretization [9, 5, 7]. However, Deep Ensembles [14] with only three models still require the most training time.

4. CONCLUSION

In this paper, we summarize the CAR MDE methods in detail along three key components, and conduct experiments on their portability and performance, including both depth and uncertainty quality. In the future, we will try to apply CAR strategy and E-Dist uncertainty estimator on more tasks.

5. REFERENCES

- [1] Gil Levi and Tal Hassner, “Age and gender classification using convolutional neural networks,” in *CVPR workshops*, 2015.
- [2] Pawan Sinha and Edward Adelson, “Recovering reflectance and illumination in a world of painted polyhedra,” in *ICCV*, 1993.
- [3] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger, “Sparsity invariant cnns,” in *3DV*, 2017.
- [4] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus, “Indoor segmentation and support inference from rgb-d images,” in *ECCV*, 2012.
- [5] Yuanzhouhan Cao, Zifeng Wu, and Chunhua Shen, “Estimating depth from monocular images as classification using deep fully convolutional residual networks,” *TCSVT*, 2017.
- [6] Bo Li, Yuchao Dai, and Mingyi He, “Monocular depth estimation with hierarchical fusion of dilated cnns and soft-weighted-sum inference,” *Pattern Recognition*, 2018.
- [7] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao, “Deep ordinal regression network for monocular depth estimation,” in *CVPR*, 2018.
- [8] Raul Diaz and Amit Marathe, “Soft labels for ordinal regression,” in *CVPR*, 2019.
- [9] Gengshan Yang, Peiyun Hu, and Deva Ramanan, “Inferring distributions over depth from a single image,” in *IROS*, 2019.
- [10] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka, “Adabins: Depth estimation using adaptive bins,” in *CVPR*, 2021.
- [11] Alex Kendall and Yarin Gal, “What uncertainties do we need in bayesian deep learning for computer vision?,” *NIPS*, 2017.
- [12] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia, “On the uncertainty of self-supervised monocular depth estimation,” in *CVPR*, 2020.
- [13] Xuanlong Yu, Gianni Franchi, and Emanuel Aldea, “Slurp: Side learning uncertainty for regression problems,” in *BMVC*, 2021.
- [14] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *NIPS*, 2017.
- [15] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan, “Enforcing geometric constraints of virtual normal for depth prediction,” in *ICCV*, 2019.
- [16] Ruibo Li, Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, and Lingxiao Hang, “Deep attention-based classification network for robust depth prediction,” in *ACCV*, 2018.
- [17] Lukas Liebel and Marco Körner, “Multidepth: Single-image depth estimation via multi-task regression and classification,” in *ITSC*, 2019.
- [18] Haoyu Ren, Mostafa El-Khamy, and Jungwon Lee, “Deep robust single image depth estimation neural network using scene understanding,” in *CVPR Workshops*, 2019.
- [19] Chen-Chou Lo and Patrick Vandewalle, “Depth estimation from monocular images and sparse radar using deep ordinal regression network,” in *ICIP*, 2021.
- [20] Minh Hieu Phan, Son Lam Phung, and Abdesselem Bouzerdoum, “Ordinal depth classification using region-based self-attention,” in *ICPR*, 2021.
- [21] Jia-Ren Chang and Yong-Sheng Chen, “Pyramid stereo matching network,” in *CVPR*, 2018.
- [22] David Eigen, Christian Puhrsch, and Rob Fergus, “Depth map prediction from a single image using a multi-scale deep network,” in *NIPS*, 2014.
- [23] Yarin Gal and Zoubin Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *ICML*, 2016.
- [24] Claude Elwood Shannon, “A mathematical theory of communication,” *ACM SIGMOBILE MC2R*, 2001.
- [25] 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.
- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell, “Fully convolutional networks for semantic segmentation,” in *CVPR*, 2015.
- [27] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger, “Densely connected convolutional networks,” in *CVPR*, 2017.
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016.
- [29] Gianni Franchi, Andrei Bursuc, Emanuel Aldea, Severine Dubuisson, and Isabelle Bloch, “One versus all for deep neural network incertitude (ovnni) quantification,” *IEEE Access*, 2021.

On Monocular Depth Estimation and Uncertainty Quantification using Classification Approaches for Regression — Supplementary Material —

Overview

Notations

Table. 4 lists some of the notations we denote and use in the Sec. 2 in the main paper.

Types	Notations	Meanings
sub/super- scripts	i	subscript for image/depth index
	j	subscript for pixel index
	p	subscript for channel index
	[method name]	superscript for indicating different methods
common capital letters	N	number of pixels with valid ground truth
	K	number of bins, the level of discretization
some other notations	$\mathbf{d}_i = \{d_{i,j}\}_j^N$	ground truth depth values
	$\hat{\mathbf{d}} = \{\hat{d}_p\}_p^K$	handcrafted logarithm depth table
	$\hat{\mathbf{d}}_i = \{\hat{d}_p\}_p^K$	learned (adaptive) depth table for a given image
	q	width between two side-by-side bins in the depth table
	$f_{\theta_1}(\mathbf{x}_i)$	output logits of the DNN f_{θ_1}
	$\hat{\mathbf{y}}_i = \{\{\hat{y}_{i,j,p}\}_p^K\}_j^N$	softmax output of the DNN f_{θ_1}

Table 4: Reminder for the notations

Original settings for CAR strategies on KITTI experiment

Table. 5 lists the original model backbone choices and experiment settings for different CAR strategies on KITTI experiment. For the coefficient used for label smoothing in Cao et al. [5], according to our discussions in the end of Sec. 2.1.2, we can transfer the original coefficient (0.5) to 65 as we report in the table using $\gamma \cdot q^{-2} = 0.5 * (\log 80/50)^{-2} = 65$.

As we can see, the previous works collected in Table. 1 in the main paper lack a full comparison, and the network structures and hyperparameters they use are also different.

Complementary to the loss function descriptions

Multiple BCE loss: is another solution when the target is a soft discrete distribution. Yang et al. [9] apply BCE loss on every class value in $\mathbf{y}_{i,j}^{[\text{smo1}]}$ defined in Eq. 5. The loss function is similar to Eq. 9:

$$L_{i,j}^{[\text{MBCE}]}(\theta_1) = -[\mathbf{y}_{i,j}^{[\text{smo1}]}] \log \delta([f_{\theta_1}(\mathbf{x}_i)]_j) + (1 - \mathbf{y}_{i,j}^{[\text{smo1}]}) \log(1 - \delta([f_{\theta_1}(\mathbf{x}_i)]_j)) \quad (15)$$

Regression loss (Smooth L1 loss): DS-SIDENet [18] applies CAR with a smooth L1 loss [21] to fit the one-hot classifica-

tion map target:

$$L_{i,j}^{[\text{smoL1}]}(\theta_1) = \begin{cases} 0.5(\hat{d}_{i,j} - k)^2 & \text{if } |\hat{d}_{i,j} - k| < 1 \\ |\hat{d}_{i,j} - k| - 0.5 & \text{otherwise} \end{cases} \quad (16)$$

$$\text{with } \hat{d}_{i,j} = \sum_{p=0}^{K-1} \hat{y}_{i,j,p} \cdot (p+1)$$

where k is defined in Eq. 3.

Regression loss (Scale-Invariant loss): is applied along with a post-processing to produce the continuous depth. Adabins [10] uses the per-image adaptive depth table $\hat{\mathbf{d}}_i^{[\text{ada}]}$ defined in Eq. 8 instead of the fixed $\hat{\mathbf{d}}$, then applies a Scale-Invariant loss [22]:

$$L_i^{[\text{SI}]}(\theta_1, \theta_2) = \omega \sqrt{\frac{1}{N} \sum_{j=0}^N h_{i,j}^2 - \frac{\lambda}{N^2} \left(\sum_{j=0}^N h_{i,j} \right)^2} \quad (17)$$

$$\text{with } h_{i,j} = \log \hat{d}_{i,j}^{[\text{ada}]} - \log d_{i,j}; \quad \hat{d}_{i,j}^{[\text{ada}]} = \sum_{p=0}^{K-1} \hat{y}_{i,j,p}^{[\text{ada}]} \cdot \hat{d}_{i,p}^{[\text{ada}]}$$

where ω and λ are hyper-parameters.

CAR MDE solutions	Encoder Backbone	Choice of K	Param. smooth	Uncertainty eval.?	Comp. with other CAR?	Comp. with Reg?
DORN [7]	ResNet-101	80	-	✗	✓	✓
Cao et al. [5]	ResNet-101	50	65	✗	✗	✓
Li et al. [6]	ResNet-152	50 (150 [16])	-	✗	✗	✓
SORN [8]	Xception	120	1	✗	✗	✓
Yang et al. [9]	ResNet-50	128	15	✓	✓	✓
DS-SIDE [18]	Self-made	80	-	✗	✓	✓
Adabins [10]	EfficientNet	256	-	✗	✓	✓

Table 5: Summary on KITTI experiment settings for typical CAR strategies. **Param. smooth:** Coefficient used for label smoothing. **Uncertainty eval.:** Whether this work evaluates the uncertainty and compared with the other works. **Comp. with other CAR:** Whether this work compared their proposed CAR strategy with the other CAR methods. **Comp. with Reg.:** Whether this work compared the CAR with the regression version of its model.

Experiments

Evaluation matrices

For the depth evaluation, we use the same matrices first introduced in [22] and used in many subsequent works. We list them as follows:

1. RMSE: $\sqrt{\frac{1}{|D|} \sum_{\mathbf{d}_i \in |D|} \|\hat{\mathbf{d}}_i - \mathbf{d}_i\|^2}$; 2. Absrel: $\frac{1}{|D|} \sum_{\mathbf{d}_i \in D} |\hat{\mathbf{d}}_i - \mathbf{d}_i| / \mathbf{d}_i$; 3. Threshold dk: Inlier metrics, k in dk indicates the power of the threshold (t), we take $t = 1.25$ and $dk =$

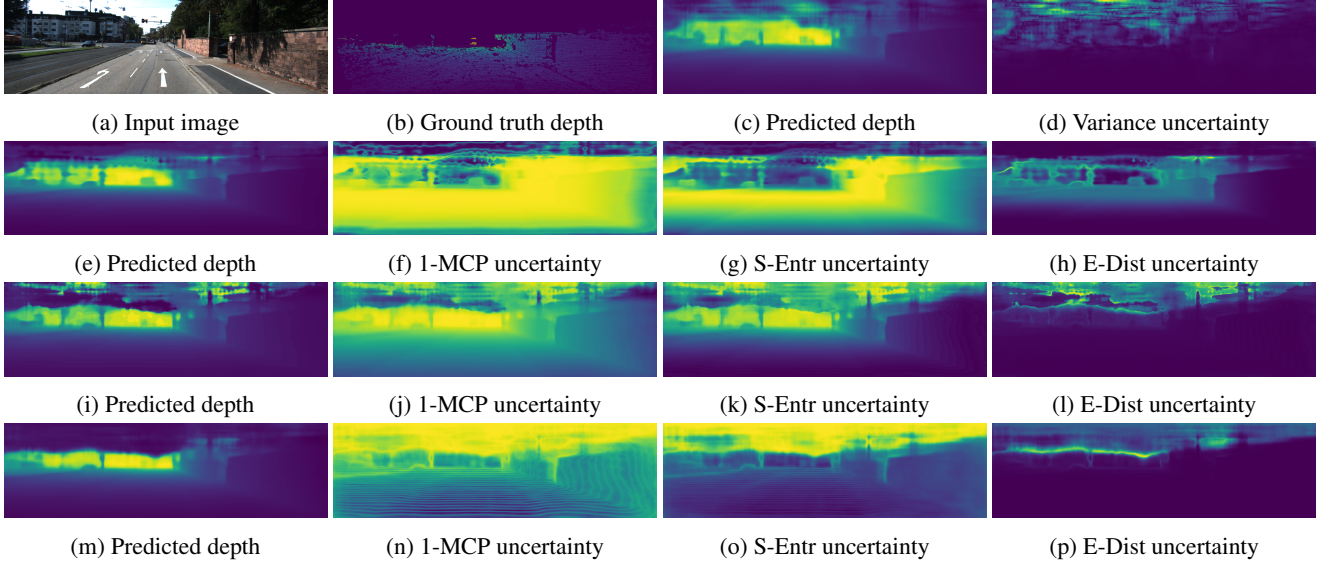


Fig. 2: Illustrations of predicted depth and uncertainty for the selected strategies applied on KITTI dataset. For both depth and the uncertainty, the brighter the pixel is, the higher depth/uncertainty value is. The figures are arranged as follows: **Input image and ground truth depth:** (a) (b). For the different strategies: **Deep Ensembles** [14]: (c) (d); **Adabins based** [10]: (e) - (h); **Dorn based** [7]: (i) - (l); **Yang et al. based** [9]: (m) - (p). All these strategies are based on FCN-ResNet101 [26, 28] backbone. The different uncertainty outputs are given by the solutions introduced in Sec. 2.2 in the main paper.

$$\frac{|A|}{|D|}, \text{ where } A = \left\{ \mathbf{x}_i, \text{ such that } \delta_i = \max\left(\frac{\hat{\mathbf{d}}_i}{\mathbf{d}_i}, \frac{\mathbf{d}_i}{\hat{\mathbf{d}}_i}\right) \text{ and } \delta_i < t^k \right\};$$

4. SqRel: $\frac{1}{|D|} \sum_{\mathbf{d}_i \in D} \|\hat{\mathbf{d}}_i - \mathbf{d}_i\|^2 / \mathbf{d}_i$;

5. RMSElog: $\sqrt{\frac{1}{|D|} \sum_{\mathbf{d}_i \in |D|} \|\log \hat{\mathbf{d}}_i - \log \mathbf{d}_i\|^2}$; 6. log10: $\frac{1}{|D|} \sum_{\mathbf{d}_i \in |D|} |\log_{10} \hat{\mathbf{d}}_i - \log_{10} \mathbf{d}_i|$.

Illustrations

Fig. 2 shows some illustrations for predicted depth as well as the predicted uncertainty given by different uncertainty estimation strategies.