Published in final edited form as:

Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit. 2018 June ; 2018: 2002–2011. doi:10.1109/CVPR.2018.00214.

Deep Ordinal Regression Network for Monocular Depth Estimation

Huan Fu¹, Mingming Gong^{2,3}, Chaohui Wang⁴, Kayhan Batmanghelich², and Dacheng Tao¹

Huan Fu: hufu6371@sydney.edu.au; Mingming Gong: mig73@pitt.edu; Chaohui Wang: chaohui.wang@u-pem.fr; Kayhan Batmanghelich: kayhan@pitt.edu; Dacheng Tao: dacheng.tao@sydney.edu.au

¹UBTECH Sydney AI Centre, SIT, FEIT, The University of Sydney, Australia

²Department of Biomedical Informatics, University of Pittsburgh

³Department of Philosophy, Carnegie Mellon University

⁴Université Paris-Est, LIGM (UMR 8049), CNRS, ENPC, ESIEE Paris, UPEM, Marne-la-Vallée, France

Abstract

Monocular depth estimation, which plays a crucial role in understanding 3D scene geometry, is an ill-posed problem. Recent methods have gained significant improvement by exploring image-level information and hierarchical features from deep convolutional neural networks (DCNNs). These methods model depth estimation as a regression problem and train the regression networks by minimizing mean squared error, which suffers from slow convergence and unsatisfactory local solutions. Besides, existing depth estimation networks employ repeated spatial pooling operations, resulting in undesirable low-resolution feature maps. To obtain high-resolution depth maps, skip-connections or multilayer deconvolution networks are required, which complicates network training and consumes much more computations. To eliminate or at least largely reduce these problems, we introduce a spacing-increasing discretization (SID) strategy to discretize depth and recast depth network learning as an ordinal regression problem. By training the network using an ordinary regression loss, our method achieves much higher accuracy and faster convergence in synch. Furthermore, we adopt a multi-scale network structure which avoids unnecessary spatial pooling and captures multi-scale information in parallel. The proposed deep ordinal regression network (DORN) achieves state-of-the-art results on three challenging benchmarks, i.e., KITTI [16], Make3D [49], and NYU Depth v2 [41], and outperforms existing methods by a large margin.

1. Introduction

Estimating depth from 2D images is a crucial step of scene reconstruction and understanding tasks, such as 3D object recognition, segmentation, and detection. In this paper, we examine the problem of *Monocular Depth Estimation* from a single image (abbr. as *MDE* hereafter).

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

Compared to depth estimation from stereo images or video sequences, in which significant progresses have been made [19, 29, 26, 44], the progress of MDE is slow. MDE is an ill-posed problem: a single 2D image may be produced from an infinite number of distinct 3D scenes. To overcome this inherent ambiguity, typical methods resort to exploiting statistically meaningful monocular cues or features, such as perspective and texture information, object sizes, object locations, and occlusions [49, 24, 32, 48, 26].

Recently, some works have significantly improved the MDE performance with the use of DCNN-based models [38, 55, 46, 9, 28, 31, 33, 3], demonstrating that deep features are superior to handcrafted features. These methods address the MDE problem by learning a DCNN to estimate the continuous depth map. Since this problem is a standard regression problem, mean squared error (MSE) in log-space or its variants are usually adopted as the loss function. Although optimizing a regression network can achieve a reasonable solution, we find that the convergence is rather slow and the final solution is far from satisfactory.

In addition, existing depth estimation networks [9, 15, 31, 33, 38, 57] usually apply standard DCNNs designed initially for image classification in a full convolutional manner as the feature extractors. In these networks, repeated spatial pooling quickly reduce the spatial resolution of feature maps (usually stride of 32), which is undesirable for depth estimation. Though high-resolution depth maps can be obtained by incorporating higher-resolution feature maps via multi-layer deconvolutional networks [33, 15, 31], multi-scale networks [38, 9] or skip-connection [57], such a processing would not only require additional computational and memory costs, but also complicate the network architecture and the training procedure.

In contrast to existing developments for MDE, we propose to discretize continuous depth into a number of intervals and cast the depth network learning as an ordinal regression problem, and present how to involve ordinal regression into a dense prediction task via DCNNs. More specifically, we propose to perform the discretization using a spacing-increasing discretization (SID) strategy instead of the uniform discretization (UD) strategy, motivated by the fact that the uncertainty in depth prediction increases along with the underlying ground-truth depth, which indicates that it would be better to allow a relatively larger error when predicting a larger depth value to avoid over-strengthened influence of large depth values on the training process. After obtaining the discrete depth values, we train the network by an ordinal regression loss, which takes into account the ordering of discrete depth values.

To ease network training and save computational cost, we introduce a network architecture which avoids unnecessary subsampling and captures multi-scale information in a simpler way instead of skip-connections. Inspired by recent advances in scene parsing [60, 4, 62], we first remove subsampling in the last few pooling layers and apply dilated convolutions to obtain large receptive fields. Then, multi-scale information is extracted from the last pooling layer by applying dilated convolution with multiple dilation rates. Finally, we develop a full-image encoder which captures image-level information efficiently at a significantly lower cost of memory than the fully-connected full-image encoders [2, 10, 9, 35, 28]. The whole

network is trained in an end-to-end manner without stage-wise training or iterative refinement. Experiments on three challenging benchmarks, *i.e.*, KITTI [16], Make3D [49, 48] and NYU Depth v2 [41], demonstrate that the proposed method achieves state-of-the-art results, and outperforms recent algorithms by a significant margin.

The remainder of this paper is organized as follows. After a brief review of related literatures in Sec. 2, we present in Sec. 3 the proposed method in detail. In Sec. 4, besides the qualitative and quantitative performance on those benchmarks, we also evaluate multiple basic instantiations of the proposed method to analyze the effects of those core factors. Finally, we conclude the whole paper in Sec. 5.

2. Related Work

Depth Estimation is essential for understanding the 3D structure of scenes from 2D images. Early works focused on depth estimation from stereo images by developing geometry-based algorithms [50, 12, 11] that rely on point correspondences between images and triangulation to estimate the depth. In a seminal work [48], Saxena *et al.* learned the depth from monocular cues in 2D images via supervised learning. Since then, a variety of approaches have been proposed to exploit the monocular cues using hand-crafted representations [49, 24, 32, 36, 7, 30, 1, 53, 45, 14, 20, 59]. Since handcrafted features alone can only capture local information, probabilistic graphic models such as Markov Random Fields (MRFs) are often built based on these features to incorporate long-range and global cues [49, 63, 39]. Another successful way to make use of global cues is the *DepthTransfer* method [26] which uses GIST global scene features [43] to search for candidate images that are “similar” to the input image from a database containing RGBD images.

Given the success of DCNNs in image understanding, many depth estimation networks have been proposed in recent years [18, 61, 35, 40, 52, 56, 46, 38, 27]. Thanks to multi-level contextual and structural information from powerful very deep networks (*e.g.*, *VGG* [54] and *ResNet* [22]), depth estimation has been boosted to a new accuracy level [9, 15, 31, 33, 57]. The main hurdle is that the repeated pooling operations in these deep feature extractors quickly decrease the spatial resolution of feature maps (usually stride 32). Eigen *et al.* [10, 9] applied multi-scale networks which stage-wisely refine estimated depth map from low spatial resolution to high spatial resolution via independent networks. Xie *et al.* [57] adopted the skip-connection strategy to fuse low-spatial resolution depth map in deeper layers with high-spatial resolution depth map in lower layers. More recent works [15, 31, 33] apply multi-layer deconvolutional networks to recover coarse-to-fine depth. Rather than solely relying on deep networks, some methods incorporate conditional random fields to further improve the quality of estimated depth maps [55, 38]. To improve efficiency, Roy and Todorovic [46] proposed the Neural Regression Forest method which allows for parallelizable training of “shallow” CNNs.

Recently, unsupervised or semi-supervised learning is introduced to learn depth estimation networks [15, 31]. These methods design reconstruction losses to estimate the disparity map by recovering a right view with a left view. Also, some weakly-supervised methods

considering pair-wise ranking information were proposed to roughly estimate and compare depth [64, 6].

Ordinal Regression [23, 21] aims to learn a rule to predict labels from an ordinal scale. Most literatures modify well-studied classification algorithms to address ordinal regression algorithms. For example, Shashua and Levin [51] handled multiple thresholds by developing a new SVM. Cammer and Singer [8] generalized the online perceptron algorithms with multiple thresholds to do ordinal regression. Another way is to formulate ordinal regression as a set of binary classification subproblems. For instance, Frank and Hall [13] applied some decision trees as binary classifiers for ordinal regression. In computer vision, ordinal regression has been combined with DCNNs to address the age estimation problem [42].

3. Method

This section first introduces the architecture of our deep ordinal regression network; then presents the SID strategy to divide continuous depth values into discrete values; and finally details how the network parameters can be learned in the ordinal regression framework.

3.1. Network Architecture

As shown in Fig. 2, the divided network consists of two parts, *i.e.*, a dense feature extractor and a scene understanding modular, and outputs multi-channel dense ordinal labels given an image.

3.1.1 Dense Feature Extractor—Previous depth estimation networks [9, 15, 31, 33, 38, 57] usually apply standard DCNNs originally designed for image recognition as the feature extractor. However, the repeated combination of max-pooling and striding significantly reduces the spatial resolution of the feature maps. Also, to incorporate multi-scale information and reconstruct high-resolution depth maps, some partial remedies, including stage-wise refinement [10, 9], skip connection [57] and multi-layer deconvolution network [15, 31, 33] can be adopted, which nevertheless not only requires additional computational and memory cost, but also complicates the network architecture and the training procedure. Following some recent scene parsing network [60, 4, 62], we advocate removing the last few downsampling operators of DCNNs and inserting holes to filters in the subsequent *conv* layers, called dilated convolution, to enlarge the field-of-view of filters without decreasing spatial resolution or increasing number of parameters.

3.1.2 Scene Understanding Modular—The scene understanding modular consists of three parallel components, *i.e.*, an atrous spatial pyramid pooling (*ASPP*) module [5], a cross-channel learner, and a full-image encoder. *ASPP* is employed to extract features from multiple large receptive fields via dilated convolutional operations. The dilation rates are 6, 12 and 18, respectively. The pure 1×1 convolutional branch can learn complex cross-channel interactions. The full-image encoder captures global contextual information and can greatly clarify local confusions in depth estimation [55, 10, 9, 2].

Though previous methods have incorporated full-image encoders, our full-image encoder contains fewer parameters. As shown in Fig. 3, to obtain global feature $\mathcal{F}$ with dimension $\mathcal{C} \times h \times w$ from F with dimension $C \times h \times w$, a common *fc*-fashion method accomplishes this by using fully-connected layers, where each element in $\mathcal{F}$ connects to all the image features, implying a global understanding of the entire image. However, this method contains a prohibitively large number of parameters, which is difficult to train and is memory consuming. In contrast, we first make use of an average pooling layer with a small kernel size and stride to reduce the spatial dimensions, followed by a *fc* layer to obtain a feature vector with dimension $\mathcal{C}$. Then, we treat the feature vector as $\mathcal{C}$ channels of feature maps with spatial dimensions of 1×1 , and add a *conv* layer with the kernel size of 1×1 as a cross-channel parametric pooling structure. Finally, we copy the feature vector to $\mathcal{F}$ along spatial dimensions so that each location of $\mathcal{F}$ share the same understanding of the entire image.

The obtained features from the aforementioned components are concatenated to achieve a comprehensive understanding of the input image. Also, we add two additional convolutional layers with the kernel size of 1×1 , where the former one reduces the feature dimension and learns complex cross-channel interactions, and the later one transforms the features into multi-channel dense ordinal labels.

3.2. Spacing-Increasing Discretization

To quantize a depth interval $[\alpha, \beta]$ into a set of representative discrete values, a common way is the uniform discretization (UD). However, as the depth value becomes larger, the information for depth estimation is less rich, meaning that the estimation error of larger depth values is generally larger. Hence, using the UD strategy would induce an over-strengthened loss for the large depth values. To this end, we propose to perform the discretization using the SID strategy (as shown in Fig. 4), which uniformed discretizes a given depth interval in *log* space to down-weight the training losses in regions with large depth values, so that our depth estimation network is capable to more accurately predict relatively small and medium depth and to rationally estimate large depth values. Assuming that a depth interval $[\alpha, \beta]$ needs to be discretized into K sub-intervals, UD and SID can be formulated as:

$$\begin{aligned} \text{UD: } t_i &= \alpha + (\beta - \alpha) * i / K, \\ \text{SID: } t_i &= e^{\log(\alpha) + \frac{\log(\beta/\alpha) * i}{K}}, \end{aligned} \quad (1)$$

where $t_i \in \{t_0, t_1, \dots, t_K\}$ are discretization thresholds. In our paper, we add a shift ξ to both α and β to obtain α^* and β^* so that $\alpha^* = \alpha + \xi = 1.0$, and apply SID on $[\alpha^*, \beta^*]$.

3.3. Learning and Inference

After obtaining the discrete depth values, it is straightforward to turn the standard regression problem into a multi-class classification problem, and adopts *softmax* regression loss to learn the parameters in our depth estimation network. However, typical multi-class

classification losses ignore the ordered information between the discrete labels, while depth values have a strong ordinal correlation since they form a well-ordered set. Thus, we cast the depth estimation problem as an ordinal regression problem and develop an ordinal loss to learn our network parameters.

Let $\chi = \phi(I, \Phi)$ denote the feature maps of size $W \times H \times C$ given an image I , where ϕ is the parameters involved in the dense feature extractor and the scene understanding modular. $Y = \psi(\chi, \Theta)$ of size $W \times H \times 2K$ denotes the ordinal outputs for each spatial locations, where $\Theta = (\theta_0, \theta_1, \dots, \theta_{2K-1})$ contains weight vectors. And $l_{(w,h)} \in \{0, 1, \dots, K-1\}$ is the discrete label produced by SID at spatial location (w, h) . Our ordinal loss $\mathcal{L}(\chi, \Theta)$ is defined as the average of pixelwise ordinal loss $\Psi(h, w, \chi, \Theta)$ over the entire image domain:

$$\mathcal{L}(\chi, \Theta) = -\frac{1}{\mathcal{N}} \sum_{w=0}^{W-1} \sum_{h=0}^{H-1} \Psi(w, h, \chi, \Theta),$$

$$\Psi(h, w, \chi, \Theta) = \sum_{k=0}^{l_{(w,h)}-1} \log(\mathcal{P}_{(w,h)}^k) + \sum_{k=l_{(w,h)}}^{K-1} (1 - \log(\mathcal{P}_{(w,h)}^k)), \quad (2)$$

$$\mathcal{P}_{(w,h)}^k = P(\hat{l}_{(w,h)} > k | \chi, \Theta),$$

where $\mathcal{N} = W \times H$, and $\hat{l}_{(w,h)}$ is the estimated discrete value decoding from $y_{(w,h)}$. We choose *softmax* function to compute $\mathcal{P}_{(w,h)}^k$ from $y_{(w,h,2k)}$ and $y_{(w,h,2k+1)}$ as follows:

$$\mathcal{P}_{(w,h)}^k = \frac{e^{y_{(w,h,2k+1)}}}{e^{y_{(w,h,2k)}} + e^{y_{(w,h,2k+1)}}}, \quad (3)$$

where $y_{(w,h,i)} = \theta_i^T x_{(w,h)}$, and $x_{(w,h)} \in \chi$. Minimizing $\mathcal{L}(\chi, \Theta)$ ensures that predictions farther from the true label incur a greater penalty than those closer to the true label.

The minimization of $\mathcal{L}(\chi, \Theta)$ can be done via an iterative optimization algorithm. Taking derivate with respect to θ_p , the gradient takes the following form:

$$\begin{aligned}
\frac{\partial \mathcal{L}(\chi, \Theta)}{\partial \theta_i} &= -\frac{1}{\mathcal{N}} \sum_{w=0}^{W-1} \sum_{h=0}^{H-1} \frac{\partial \Psi(w, h, \chi, \Theta)}{\partial \theta_i}, \\
\frac{\partial \Psi(w, h, \chi, \Theta)}{\partial \theta_{2k+1}} &= -\frac{\partial \Psi(w, h, \chi, \Theta)}{\partial \theta_{2k}}, \\
\frac{\partial \Psi(w, h, \chi, \Theta)}{\partial \theta_{2k}} &= x_{(w,h)} \eta(l_{(w,h)} > k) (\mathcal{P}_{(w,h)}^k - 1) + x_{(w,h)} \eta(l_{(w,h)} \leq k) \mathcal{P}_{(w,h)}^k,
\end{aligned} \tag{4}$$

an where $k \in \{0, 1, \dots, K-1\}$, and $\eta(\cdot)$ is indicator function such that $\eta(\text{true}) = 1$ and $\eta(\text{false}) = 0$. We the can optimize our network via backpropagation.

In the inference phase, after obtaining ordinal labels for each position of image I , the predicted depth value $\hat{d}_{(w,h)}$ is decoded as:

$$\begin{aligned}
\hat{d}_{(w,h)} &= \frac{t_{\hat{l}_{(w,h)}} + t_{\hat{l}_{(w,h)} + 1}}{2} - \xi, \\
\hat{l}_{(w,h)} &= \sum_{k=0}^{K-1} \eta(\mathcal{P}_{(w,h)}^k > 0.5).
\end{aligned} \tag{5}$$

4. Experiments

To demonstrate the effectiveness of our depth estimator, we present a number of experiments examining different aspects of our approach. After introducing the implementation details, we evaluate our methods on three challenging outdoor datasets, *i.e.* *KITTI* [16], *Make3D* [48, 49] and *NYU Depth v2* [41]. The evaluation metrics are following previous works [10, 38]. Some ablation studies based on *KITTI* are discussed to give a more detailed analysis of our method.

Implementation Details We implement our depth estimation network based on the public deep learning platform *Caffe* [25]. The learning strategy applies a polynomial decay with a base learning rate of 0.0001 and the power of 0.9. Momentum and weight decay are set to 0.9 and 0.0005 respectively. The iteration number is set to 300K for *KITTI*, 50K for *Make3D*, and 3M for *NYU Depth v2*, and batch size is set to 3. We find that further increasing the iteration number can only slightly improve the performance. We adopt both *VGG-16* [54] and *ResNet-101* [22] as our feature extractors, and initialize their parameters via the pre-trained classification model on ILSVRC [47]. Since features in first few layers only contain general low-level information, we fixed the parameters of *conv1* and *conv2* blocks in *ResNet* after initialization. Also, the batch normalization parameters in *ResNet* are directly initialized and fixed during training progress. Data augmentation strategies are following [10]. In the test phase, we split each image to some overlapping windows according the cropping method in the training phase, and obtain the predicted depth values in overlapped regions by averaging the predictions.

4.1. Benchmark Performance

KITTI The KITTI dataset [16] contains outdoor scenes with images of resolution about 375×1241 captured by cameras and depth sensors in a driving car. All the 61 scenes from the “city”, “residential”, “road” and “Campus” categories are used as our training/test sets. We test on 697 images from 29 scenes split by Eigen *et al.* [10], and train on about 23488 images from the remaining 32 scenes. We train our model on a random crop of size 385×513 . For some other details, we set the maximal ordinal label for KITTI as 80, and evaluate our results on a pre-defined center cropping following [10] with the depth ranging from $0m$ to $80m$ and $0m$ to $50m$. Note that, a single model is trained on the full depth range, and is tested on data with different depth ranges.

Make3D The Make3D dataset [48, 49] contains 534 out-door images, 400 for training, and 134 for testing, with the resolution of 2272×1704 , and provides the ground truth depth map with a small resolution of 55×305 . We reduce the resolution of all images to 568×426 , and train our model on a random crop of size 513×385 . Following previous works, we report $C1$ (depth range from $0m$ to $80m$) and $C2$ (depth range from $0m$ to $70m$) error on this dataset using three commonly used evaluation metrics [26, 38]. For the VGG model, we train our DORN on a depth range of $0m$ to $80m$ from scratch (ImageNet model), and evaluate results using the same model for $C1$ and $C2$. However, for ResNet, we learn two separate models for $C1$ and $C2$ respectively.

NYU Depth v2 The NYU Depth v2 [41] dataset contains 464 indoor video scenes taken with a Microsoft Kinect camera. We train our DORN using all images (about 120K) from the 249 training scenes, and test on the 694-image test set following previous works. To speed up training, all the images are reduced to the resolution of 288×384 from 480×640 . And the model are trained on random crops of size 257×353 . We report our scores on a pre-defined center cropping by Eigen [10].

Performance Tab. 2 and Tab. 3 give the results on two out-door datasets, i.e., KITTI and Make3D. It can be seen that our DORN improves the accuracy by 5% ~ 30% in terms of all metrics compared with previous works in all settings. Some qualitative results are shown in Fig. 5 and Fig. 6. In Tab. 4, our DORN outperforms other methods on NYU Depth v2, which is one of the largest indoor benchmarks. The results suggest that our method is applicable to both indoor and outdoor data. We also evaluate our method on the online KITTI evaluation server¹. As shown in Tab. 1, our DORN significantly outperforms the officially provided baseline about 30% ~ 70%.

4.2. Ablation Studies

We conduct various ablation studies to analyze the details of our approach. Results are shown in Tab. 5, Tab. 6, Fig. 1, and Fig. 7, and discussed in detail.

4.2.1 Depth Discretization—Depth discretization is critical to performance improvement, because it allows us to apply classification and ordinal regression losses to optimize the network parameters. According to scores in Tab. 5, training by regression on continuous depth seems to converge to a poorer solution than the other two methods, and our

ordinal regression network achieves the best performance. There is an obvious gap between approaches where depth is discretized by SID and UD, respectively. Besides, when replacing our ordinal regression loss by an advantage regression loss (*i.e.* BerHu), our DORN still obtain much higher scores. Thus, we can conclude that: (i) SID is important and can further improve the performance compared to UD; (ii) discretizing depth and training using a multi-class classification loss is better than training using regression losses; (iii) exploring the ordinal correlation among depth drives depth estimation networks to converge to even better solutions.

Furthermore, we also train the network using $\text{RMSE}_{\log}$ on discrete depth values obtained by SID, and report the results in Tab. 5. We can see that MSE-SID performs slightly better than MSE, which demonstrates that quantization errors are nearly ignorable in depth estimation. The benefits of discretization through the use of ordinal regression losses far exceeds the cost of depth discretization.

4.2.2 Full-image Encoder—From Tab. 6, a full-image encoder is important to further boost the performance. Our full-image encoder yields a little higher scores than fc type encoders [2, 10, 9, 35, 28], but significantly reduce the number of parameters. For example, we set C to 512 (VGG), $\mathcal{C}$ to 512, m to 2048 (Eigen [10, 9]), and k to 4 in Fig. 3. Because of limited computation resources, when implementing the fc -fashion encoder, we downsampled the resolution of F using the stride of 3, and upsampled $\mathcal{F}$ to the required resolution. With an input image of size 385×513 , h and w will be 49 and 65 respectively in our network. The number of parameters in fc -fashion encoder and our encoder is

$$\frac{1}{9} * m * w * h * C + m^2 + \frac{1}{9} * w * h * \mathcal{C} * m \approx 753M, \text{ and is } \mathcal{C} * \frac{w}{4} * \frac{h}{4} * C + \mathcal{C} * \mathcal{C} \approx 51M,$$

respectively. From the experimental results and parameter analysis, it can be seen that our full-image encoder performs better while requires less computational resources.

4.2.3 How Many Intervals—To illustrate the sensitivity to the number of intervals, we discretizing depth into various number of intervals via SID. As shown in Fig. 7, with a range of 40 to 120 intervals, our DORN has a score in [0.908, 0.915] regarding $\delta < 1.25$, and a score in [3.056, 3.125] in terms of RMSE, and is thereby robust to a long range of depth interval numbers. We can also see that neither too few nor too many depth intervals are rational for depth estimation: too few depth intervals cause large quantization error, while too many depth intervals lose the advantage of discretization.

5. Conclusion

In this paper, we have developed an deep ordinal regression network (DORN) for monocular depth estimation MDE from a single image, consisting of a clean CNN architecture and some effective strategies for network optimization. Our method is motivated by two aspects: (i) to obtain high-resolution depth map, previous depth estimation networks require incorporating multi-scale features as well as full-image features in a complex architecture, which complicates network training and largely increases the computational cost; (ii) training a regression network for depth estimation suffers from slow convergence and

unsatisfactory local solutions. To this end, we first introduced a simple depth estimation network which takes advantage of dilated convolution technique and a novel full-image encoder to directly obtain a high-resolution depth map. Moreover, an effective depth discretization strategy and an ordinal regression training loss were intergrated to improve the training of our network so as to largely increase the estimation accuracy. The proposed method achieves the state-of-the-art performance on the KITTI, Make3D and NYU Depth v2 datasets. In the future, we will investigate new approximations to depth and extend our framework to other dense prediction problems.

Acknowledgement

This research was supported by Australian Research Council Projects FL-170100117 and DP-180103424. This work was partially supported by SAP SE and CNRS INS2IJJC-INVISANA. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research. This research was partially supported by research grant from Pfizer titled "Developing Statistical Method to Jointly Model Genotype and High Dimensional Imaging Endophenotype." We were also grateful for the computational resources provided by Pittsburgh Super Computing grant number TG-ASC170024.

References

- [1]. Baig MH and Torresani L. Coupled depth learning In WACV, 2016.
- [2]. Chakrabarti A, Shao J, and Shakhnarovich G. Depth from a single image by harmonizing overcomplete local network predictions In NIPS, 2016.
- [3]. Chang A, Dai A, Funkhouser T, Halber M, Niessner M, Savva M, Song S, Zeng A, and Zhang Y. Matterport3D: Learning from RGB-D data in indoor environments. 3DV, 2017.
- [4]. Chen L-C, Papandreou G, Kokkinos I, Murphy K, and Yuille AL. Semantic image segmentation with deep convolutional nets and fully connected crfs In ICLR, 2015.
- [5]. Chen L-C, Papandreou G, Kokkinos I, Murphy K, and Yuille AL. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. arXiv: 1606.00915, 2016.
- [6]. Chen W, Fu Z, Yang D, and Deng J. Single-image depth perception in the wild In NIPS, 2016.
- [7]. Choi S, Min D, Ham B, Kim Y, Oh C, and Sohn K. Depth analogy: Data-driven approach for single image depth estimation using gradient samples. IEEE TIP, 24(12):5953–5966, 2015.
- [8]. Crammer K and Singer Y. Pranking with ranking In NIPS, 2002.
- [9]. Eigen D and Fergus R. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture In ICCV, 2015.
- [10]. Eigen D, Puhrsch C, and Fergus R. Depth map prediction from a single image using a multi-scale deep network In NIPS, 2014.
- [11]. Flynn J, Neulander I, Philbin J, and Snavely N. Deep-stereo: Learning to predict new views from the world's imagery In CVPR, 2016.
- [12]. Forsyth D and Ponce J. Computer Vision: a Modern Approach. Prentice Hall, 2002.
- [13]. Frank E and Hall M. A simple approach to ordinal classification. ECML, 2001.
- [14]. Furukawa R, Sagawa R, and Kawasaki H. Depth estimation using structured light flow – analysis of projected pattern flow on an object's surface In ICCV, 2017.
- [15]. Garg R, Carneiro G, and Reid I. Unsupervised cnn for single view depth estimation: Geometry to the rescue In ECCV, 2016.
- [16]. Geiger A, Lenz P, Stiller C, and Urtasun R. Vision meets robotics: The kitti dataset. IJRR, 2013.
- [17]. Godard C, Mac Aodha O, and Brostow GJ. Unsupervised monocular depth estimation with left-right consistency. CVPR, 2017.
- [18]. Güler RA, Trigeorgis G, Antonakos E, Snape P, Zafeiriou S, and Kokkinos I. Densereg: Fully convolutional dense shape regression in-the-wild In CVPR, 2016.

- [19]. Ha H, Im S, Park J, Jeon H-G, and Kweon IS. High-quality depth from uncalibrated small motion clip In CVPR, 2016.
- [20]. Hane C, Ladicky L, and Pollefeys M. Direction matters: Depth estimation with a surface normal classifier In CVPR, 2015.
- [21]. Harrell FE Jr. Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis. Springer, 2015.
- [22]. He K, Zhang X, Ren S, and Sun J. Deep residual learning for image recognition In CVPR, 2016.
- [23]. Herbrich R, Graepel T, and Obermayer K. Support vector learning for ordinal regression. 1999.
- [24]. Hoiem D, Efros AA, and Hebert M. Recovering surface layout from an image. IJCV, 75(1):151–172, 2007.
- [25]. Jia Y, Shelhamer E, Donahue J, Karayev S, Long J, Girshick R, Guadarrama S, and Darrell T. Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093, 2014.
- [26]. Karsch K, Liu C, and Kang SB. Depth transfer: Depth extraction from video using non-parametric sampling. IEEE TPAMI, 36(11):2144–2158, 2014.
- [27]. Kendall A and Gal Y. What uncertainties do we need in bayesian deep learning for computer vision? In NIPS, 2017.
- [28]. Kim S, Park K, Sohn K, and Lin S. Unified depth prediction and intrinsic image decomposition from a single image via joint convolutional neural fields In ECCV, 2016.
- [29]. Kong N and Black MJ. Intrinsic depth: Improving depth transfer with intrinsic images In ICCV, 2015.
- [30]. Konrad J, Wang M, Ishwar P, Wu C, and Mukherjee D. Learning-based, automatic 2d-to-3d image and video conversion. IEEE TIP, 22(9):3485–3496, 2013.
- [31]. Kuznetsov Y, Stückler J, and Leibe B. Semi-supervised deep learning for monocular depth map prediction CVPR, 2017.
- [32]. Ladicky L, Shi J, and Pollefeys M. Pulling things out of perspective In CVPR, 2014.
- [33]. Laina I, Rupprecht C, Belagiannis V, Tombari F, and Navab N. Deeper depth prediction with fully convolutional residual networks In 3DV, 2016.
- [34]. Li B, Shen C, Dai Y, van den Hengel A, and He M. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs In CVPR, 2015.
- [35]. Li J, Klein R, and Yao A. A two-streamed network for estimating fine-scaled depth maps from single rgb images In ICCV, 2017.
- [36]. Li X, Qin H, Wang Y, Zhang Y, and Dai Q. Dept: depth estimation by parameter transfer for single still images In ACCV, 2014.
- [37]. Liu B, Gould S, and Koller D. Single image depth estimation from predicted semantic labels In CVPR, 2010.
- [38]. Liu F, Shen C, Lin G, and Reid I. Learning depth from single monocular images using deep convolutional neural fields. IEEE TPAMI, 38(10):2024–2039, 2016.
- [39]. Liu M, Salzmann M, and He X. Discrete-continuous depth estimation from a single image In CVPR, 2014.
- [40]. Narihira T, Maire M, and Yu SX. Learning lightness from human judgement on relative reflectance In CVPR, 2015.
- [41]. K. P Silberman Nathan, Hoiem Derek and Fergus R. Indoor segmentation and support inference from rgbd images In ECCV, 2012.
- [42]. Niu Z, Zhou M, Wang L, Gao X, and Hua G. Ordinal regression with multiple output cnn for age estimation In CVPR, 2016.
- [43]. Oliva A and Torralba A. Modeling the shape of the scene: A holistic representation of the spatial envelope. IJCV, 42(3):145–175, 2001.
- [44]. Rajagopalan A, Chaudhuri S, and Mudenagudi U. Depth estimation and image restoration using defocused stereo pairs. IEEE TPAMI, 26(11):1521–1525, 2004.
- [45]. Ranftl R, Vineet V, Chen Q, and Koltun V. Dense monocular depth estimation in complex dynamic scenes In CVPR, 2016.

- [46]. Roy A and Todorovic S. Monocular depth estimation using neural regression forest In CVPR, 2016.
- [47]. Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, Berg AC, and Fei-Fei L. ImageNet Large Scale Visual Recognition Challenge. IJCV, 115(3):211–252, 2015.
- [48]. Saxena A, Chung SH, and Ng AY. Learning depth from single monocular images In NIPS, 2006.
- [49]. Saxena A, Sun M, and Ng AY. Make3d: Learning 3d scene structure from a single still image. IEEE TPAMI, 31(5):824–840, 2009.
- [50]. Scharstein D and Szeliski R. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. IJCV, 47(1–3):7–42, 2002.
- [51]. Shashua A and Levin A. Ranking with large margin principle: Two approaches In NIPS, 2003.
- [52]. Shelhamer E, Barron JT, and Darrell T. Scene intrinsics and depth from a single image In ICCV Workshop, 2015.
- [53]. Shi J, Tao X, Xu L, and Jia J. Break ames room illusion: depth from general single images. ACM TOG, 34(6):225, 2015.
- [54]. Simonyan K and Zisserman A. Very deep convolutional networks for large-scale image recognition In ICLR, 2015.
- [55]. Wang P, Shen X, Lin Z, Cohen S, Price B, and Yuille A. Towards unified depth and semantic prediction from a single image In CVPR, 2015.
- [56]. Wang X, Fouhey D, and Gupta A. Designing deep networks for surface normal estimation In CVPR, 2015.
- [57]. Xie J, Girshick R, and Farhadi A. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks In ECCV, 2016.
- [58]. Xu D, Ricci E, Ouyang W, Wang X, and Sebe N. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation In CVPR, 2017.
- [59]. You X, Li Q, Tao D, Ou W, and Gong M. Local metric learning for exemplar-based object detection. IEEE TCSVT, 24(8):1265–1276, 2014.
- [60]. Yu F and Koltun V. Multi-scale context aggregation by dilated convolutions In ICLR, 2016.
- [61]. Zhang Z, Schwing AG, Fidler S, and Urtasun R. Monocular object instance segmentation and depth ordering with cnns In ICCV, 2015.
- [62]. Zhao H, Shi J, Qi X, Wang X, and Jia J. Pyramid scene parsing network In CVPR, 2017.
- [63]. Zhuo W, Salzmann M, He X, and Liu M. Indoor scene structure analysis for single image depth estimation In CVPR, 2015.
- [64]. Zoran D, Isola P, Krishnan D, and Freeman WT. Learning ordinal relationships for mid-level vision In ICCV, 2015.

Figure 1: Estimated Depth by DORN.

MSE: Training our network via MSE in $\log$ space, where ground truths are continuous depth values. DORN: The proposed deep ordinal regression network. Depth values in the black part are not provided by KITTI.

Figure 2: Illustration of the network architecture.

The network consists of a dense feature extractor, multi-scale feature learner (*ASPP*), cross channel information learner (the pure 1×1 convolutional branch), a full-image encoder and an ordinal regression optimizer. The *Conv* components here are all with kernel size of 1×1 . The *ASPP* module consists of 3 dilated convolutional layers with kernel size of 3×3 and dilated rate of 6, 12 and 18 respectively. The supervised information of our network is discrete depth values output by the discretization using the *SID* strategy. The whole network is optimized by our ordinal regression training loss in an end-to-end fashion.

Figure 3: Full-Image Encoders.

Top: the full-image encoder implemented by pure *fc* layers [10, 9, 2] ($\delta < 1.25: 0.910$);

Bottom: Our proposed encoder ($\delta < 1.25: 0.915$).

Figure 4: Discrete Intervals.

Illustration of UD (middle) and SID (bottom) to discretize depth interval $[\alpha, \beta]$ into five subintervals. See Eq. 1 for details.

Figure 5: Depth Prediction on KITTI.

Image, ground truth, Eigen [10], LRC [17], and our DORN. Ground truth has been interpolated for visualization. Pixels with distance $> 80m$ in LRC are masked out.

Figure 6: Depth Prediction on Make3D.

Image, ground truth, and our DORN. Pixels with distance $> 70m$ are masked out.

Figure 7: Performance Ranging Different Intervals via SID.
 Left: accuracy on $\delta < 1.25$. Right: evaluation errors on RMSE.

Table 1:
Scores on the online KITTI evaluation server.

We train our DORN using the officially provided training and validation sets.

Method	SILog	sqErrorRel	absErrorRel	iRMSE
Official Baseline	18.19	7.32	14.24	18.50
DORN	11.80	2.19	8.93	13.22

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Performance on KITTI.

Table 2:

All the methods are evaluated on the test split by Eigen *et al.* [10]. LRC (CS + K): LRC pre-train their model on Cityscapes and fine tune on KITTI.

Method	cap	higher is better			lower is better		
		$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$	Abs Rel	Squa Rel	RMSE
Make3D [49]	0 – 80 m	0.601	0.820	0.926	0.280	3.012	8.734
Eigen <i>et al.</i> [10]	0 – 80 m	0.692	0.899	0.967	0.190	1.515	7.156
Liu <i>et al.</i> [38]	0 – 80 m	0.647	0.882	0.961	0.217	1.841	6.986
LRC (CS + K) [17]	0 – 80 m	0.861	0.949	0.976	0.114	0.898	4.935
Kuznetsov <i>et al.</i> [31]	0 – 80 m	0.862	0.960	0.986	0.113	0.741	4.621
DORN (VGG)	0 – 80 m	0.915	0.980	0.993	0.081	0.376	3.056
DORN (ResNet)	0 – 80 m	0.932	0.984	0.994	0.072	0.307	2.727
Garg <i>et al.</i> [15]	0 – 50 m	0.740	0.904	0.962	0.169	1.080	5.104
LRC (CS + K) [17]	0 – 50 m	0.873	0.954	0.979	0.108	0.657	3.729
Kuznetsov <i>et al.</i> [31]	0 – 50 m	0.875	0.964	0.988	0.108	0.595	3.518
DORN (VGG)	0 – 50 m	0.920	0.982	0.994	0.079	0.324	2.517
DORN (ResNet)	0 – 50 m	0.936	0.985	0.995	0.071	0.268	2.271
							0.116

Table 3:
Performance on Make3D.

LRC-Deep3D [57] is adopting LRC [17] on Deep3D model [57].

Method	C1 error			C2 error		
	rel	log ₁₀	rms	rel	log ₁₀	rms
Make3D [49]	-	-	-	0.370	0.187	-
Liu <i>et al.</i> [37]	-	-	-	0.379	0.148	-
DepthTransfer [26]	0.355	0.127	9.20	0.361	0.148	15.10
Liu <i>et al.</i> [39]	0.335	0.137	9.49	0.338	0.134	12.60
Li <i>et al.</i> [34]	0.278	0.092	7.12	0.279	0.102	10.27
Liu <i>et al.</i> [38]	0.287	0.109	7.36	0.287	0.122	14.09
Roy <i>et al.</i> [46]	-	-	-	0.260	0.119	12.40
Laina <i>et al.</i> [33]	0.176	0.072	4.46	-	-	-
LRC-Deep3D [57]	1.000	2.527	19.11	-	-	-
LRC [17]	0.443	0.156	11.513	-	-	-
Kuznietsov <i>et al.</i> [31]	0.421	0.190	8.24	-	-	-
MS-CRF [58]	0.184	0.065	4.38	0.198	-	8.56
DORN (VGG)	0.236	0.082	7.02	0.238	0.087	10.01
DORN (ResNet)	0.157	0.062	3.97	0.162	0.067	7.32

Table 4:
Performance on NYU Depth v2.

δ_i : $\delta < 1.25^i$; ResNet based model.

Method	δ_1	δ_2	δ_3	rel	$\log_{10}$	rms
Make3D [49]	0.447	0.745	0.897	0.349	-	1.214
DepthTransfer [26]	-	-	-	0.35	0.131	1.2
Liu <i>et al.</i> [39]	-	-	-	0.335	0.127	1.06
Ladicky <i>et al.</i> [32]	0.542	0.829	0.941	-	-	-
Li <i>et al.</i> [34]	0.621	0.886	0.968	0.232	0.094	0.821
Wang <i>et al.</i> [55]	0.605	0.890	0.970	0.220	-	0.824
Roy <i>et al.</i> [46]	-	-	-	0.187	-	0.744
Liu <i>et al.</i> [38]	0.650	0.906	0.976	0.213	0.087	0.759
Eigen <i>et al.</i> [9]	0.769	0.950	0.988	0.158	-	0.641
Chakrabarti <i>et al.</i> [2]	0.806	0.958	0.987	0.149	-	0.620
Laina <i>et al.</i> [33]	0.629	0.889	0.971	0.194	0.083	0.790
Li <i>et al.</i> [35]	0.789	0.955	0.988	0.152	0.064	0.611
Laina <i>et al.</i> [33] [†]	0.811	0.953	0.988	0.127	0.055	0.573
Li <i>et al.</i> [35] [†]	0.788	0.958	0.991	0.143	0.063	0.635
MS-CRF [58] [†]	0.811	0.954	0.987	0.121	0.052	0.586
DORN[†]	0.828	0.965	0.992	0.115	0.051	0.509

Table 5:

Depth Discretization and Ordinal Regression.

MSE: mean squared error in log space. MCC: multi-class classification. DORN: proposed ordinal regression. Note that training by MSE for 1M iterations only slightly improve the performance compared with 0.5M (about 0.001 on $\delta < 1.25$). berHu: the reverse Huber loss. [†]: ResNet based model.

Variant	Iteration	higher is better		lower is better				
		$\delta < 1:25$	$\delta < 1:25^2$	$\delta < 1:25^3$	Abs Rel	Squa Rel	RMSE	RMSE _{log}
MSE	1M	0.864	0.969	0.991	0.109	0.527	3.660	0.164
MSE-SID	0.6M	0.865	0.970	0.992	0.108	0.520	3.636	0.163
MCC-UD	0.3M	0.892	0.970	0.988	0.093	0.474	3.438	0.155
MCC-SID	0.3M	0.906	0.976	0.991	0.084	0.417	3.201	0.142
DORN-UD	0.3M	0.900	0.973	0.991	0.091	0.452	3.339	0.148
DORN-SID	0.3M	0.915	0.980	0.993	0.081	0.376	3.056	0.132
berHu [†]	0.6M	0.909	0.978	0.992	0.086	0.385	3.365	0.136
DORN [†]	0.3M	0.932	0.984	0.994	0.072	0.307	2.727	0.120

Table 6:**Full-image Encoder.**

Parameters here is computed by some common settings in Eigen [10] and our DORN.

Variant	$\delta < 1:25$	Abs Rel	RMSE _{log}	Params
w/o full-image encoder	0.906	0.092	0.143	0M
<i>f</i> c-fashion	0.910	0.085	0.137	753M
our encoder	0.915	0.081	0.132	51M

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript