

Object-aware Monocular Depth Prediction with Instance Convolutions

Enis Simsar^{1,*}, Evin Pınar Örnek^{1,*}, Fabian Manhardt², Helisa Dhamo¹, Nassir Navab¹ and Federico Tombari^{1,2}

Abstract—With the advent of deep learning, estimating depth from a single RGB image has recently received a lot of attention, being capable of empowering many different applications ranging from path planning for robotics to computational cinematography. Nevertheless, while the depth maps are in their entirety fairly reliable, the estimates around object discontinuities are still far from satisfactory. This can be attributed to the fact that the convolutional operator naturally aggregates features across object discontinuities, resulting in smooth transitions rather than clear boundaries. Therefore, in order to circumvent this issue, we propose a novel convolutional operator which is explicitly tailored to avoid feature aggregation of different object parts. In particular, our method is based on estimating per-part depth values by means of super-pixels. The proposed convolutional operator, which we dub "Instance Convolution", then only considers each object part individually on the basis of the estimated super-pixels. Our evaluation with respect to the NYUv2, iBims and KITTI datasets demonstrate the advantages of Instance Convolutions over the classical convolution at estimating depth around occlusion boundaries, while producing comparable results elsewhere. Our code is available here ¹.

I. INTRODUCTION

Monocular depth prediction (MDP) is a very important field in research due to its wide range of applications in robotics and AR [1], [2], [3]. Nevertheless, predicting accurate depth from monocular input is also an inherently ill-posed problem. For the human perceptual system depth perception is a simpler task, as we heavily rely on prior knowledge from the environment. Similarly, deep learning has recently proven to be particularly suited for such problems, as the network is also capable of leveraging visual priors when making a prediction [4], [5].

With the rise of deep learning and increasing availability of appropriate and large datasets [7], [6], [8], [9], [10], depth prediction from single images has recently made a huge leap forward in terms of robustness and accuracy [11], [12], [13], [14], [15]. Yet, despite those large improvements, they still often fall short of adequate quality for specific robotics applications, such as in path planning or robotic interventions where robots need to operate in hazardous environments with low-albedo surfaces and clutter [1], [6], [16]. One of the most limiting factors is the poor quality around object edges and surfaces, which directly affects the 3D perception, thereby resulting in a robot missing the objects. The predicted depth

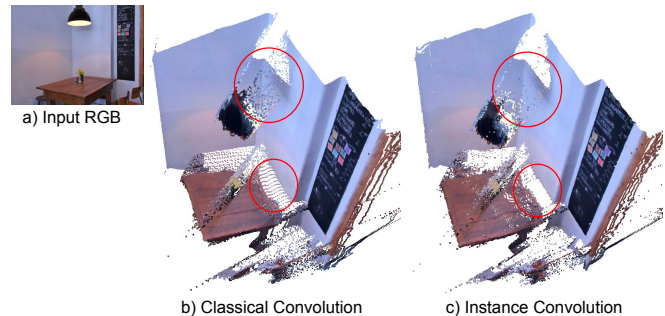


Fig. 1. The interpolation effect of classical convolutions induces smeared occlusion boundaries in the predicted depth map (here on an image taken from the iBims [6] dataset), as visible in (b), whereas the proposed Instance Convolution improves on this drawback (c). Note that, while this effect is not as evident when visualizing the 2D depth map, it gets clearly revealed once the depth map is back-projected in 3D.

maps are typically blurry around object boundaries due to the nature of 2D convolutions and bilinear upsampling. Since the kernel aggregates features across object boundaries, the estimated depth map commonly results in an undesired interpolation between fore- and background. Similarly, associated 3D point clouds cannot reflect 3D structures (see Fig. 1). In this work, our motivation is to capture object-based depth values more sharply and completely, while preserving the global consistency with the rest of the scene.

To circumvent the smeared boundary problem, *i.e.* avoid undesired depth interpolation across different segments, we are interested in an operation that extracts features within a continuous object segment. To achieve this, we employ a novel convolution operation inspired by the sparse convolutions introduced by Uhrig *et al.* [17].

Sparse convolutions are characterized by a mask, which defines the region in which the convolution operates. While sparse convolutions typically rely on a single mask throughout the entire image, in this work, the mask depends on the pixel location. Given a filter window, we define the mask as the feature region that belongs to the same segment as the central pixel of that window. In other words, only the pixels that belong to a certain object contribute to its feature extraction. We name our convolution operator *Instance Convolution*.

Using Instance Convolutions to learn the object depth values should make the depth values at the object edges sharper than regular convolution, *i.e.* prevent the interpolation problem at the occlusion boundaries. Despite this advantage in terms of boundary sharpness, Instance Convolutions come

¹:Technical University of Munich, Germany: {enis.simsar, evin.oerne, helisa.dhamo, nassir.navab, federico.tombari}@tum.de

²:Google Inc. fabianmanhardt@google.com

*:the first two authors contributed equally.

¹<https://github.com/enisimsar/instance-conv>

with an obvious drawback. An architecture based solely on such operation, would not be able to capture object extent and global context. This inherent scale-distance ambiguity would thus make it impossible to obtain metric depth. Therefore, we propose an architecture that first extracts global features so to utilize scene priors via a common backbone comprised of regular convolutions. We then append a block composed of Instance Convolutions to rectify the features within an object segment, resulting in sharper depth across occlusion boundaries. Notice that we chose an optimization-based approach [18], producing super-pixels, to obtain corresponding segmentation as a deep learning-driven method would simply shift the problem of clear boundaries towards the segmenter. In addition, our method does not require any semantic information, but rather only needs to understand what pixels belong to the same discontinuity-free object part.

Our contributions can be summarized as follows:

- We propose a novel end-to-end method for depth estimation from monocular data, which explicitly enforces clear object boundaries by means of super-pixels.
- To this end we propose a dynamic convolutional operator, *Instance Convolution*, which only aggregates features appertaining to the same segment as the center pixel, with respect to the current kernel location.
- Further, as we are required to properly propagate the correct segment information throughout the whole network, we additionally introduce the "center-pooling" operator, keeping track of center pixel's segment id.

We validate the usefulness of Instance Convolutions for edge-aware depth estimation on three commonly employed benchmarks, namely NYUv2 [7], iBims [6], and KITTI [19]. Thereby, we show that Instance Convolutions is able to improve object boundaries regardless of the chosen backbone depth estimator.

II. RELATED WORK

a) Supervised monocular depth prediction: The first attempts to tackle monocular depth estimation were proposed by [4], [20] via hand-engineered features and Markov Random Fields (MRF). Later, the advancements in deep learning established a new era for depth estimation, starting with Eigen *et al.* [21]. One of the main problems in learned depth regression occurs in the decoder part. Due to the successive layers of convolution channels in neural networks, fine details of the input images are lost. There are a number of works that approach this problem in different ways. Eigen & Fergus [22] introduced multi-scale networks to make depth predictions at multiple scales. Laina *et al.* [23] built upon a ResNet architecture with improved up-sampling blocks to reduce information loss in the decoding phase. Xu *et al.* [24] proposed an approach that combines deep learning with conditional random fields (CRF), where CRFs are used to fuse the multi-scale deep features.

A line of works pursued multitask learning approaches that predict semantic or instance labels [14], depth edges, and normals [25], [26], [27] to improve depth prediction. Kendall *et al.* [28] investigated the effect of uncertainty estimation in

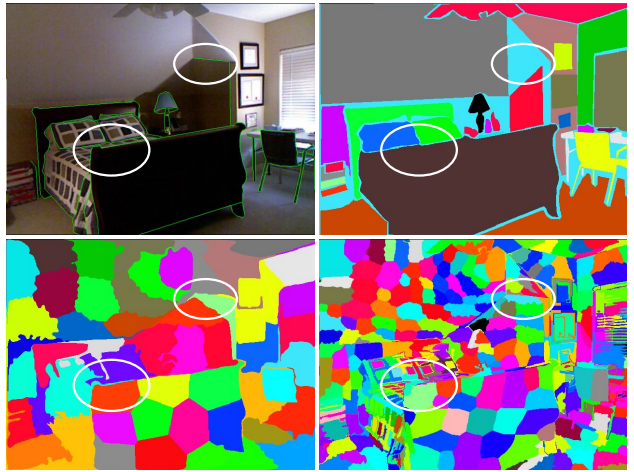


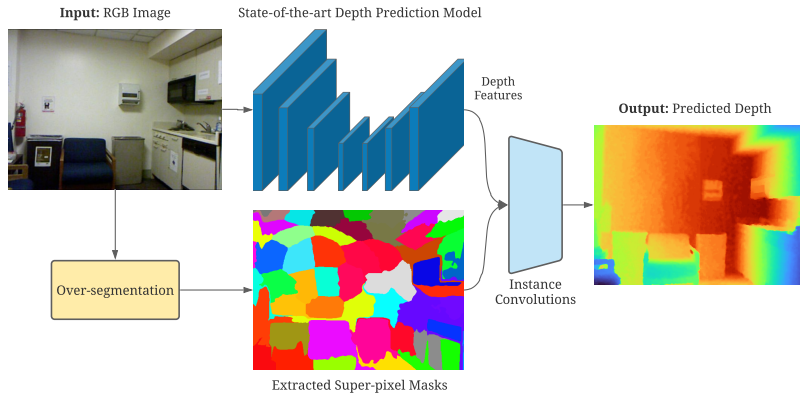
Fig. 2. **Different segmentation methods.** The top left shows the input RGB image with occlusion boundaries in green, while the top right illustrates the ground truth segmentation. Notice how the latter does not consider intra-object discontinuities (highlighted with white-circles). Thus, we leverage super-pixels to account for any discontinuities based on the RGB input. Exemplary, we demonstrate the obtained super-pixels for SLIC [18] using 64 segments on the bottom left, and BASS [40] on the bottom right. Note that BASS outputs approximately 200 segments on average.

scene understanding. In contrast, Yin *et al.* [13] estimate the 3D point cloud from a predicted depth map, using the surface geometry as additional guidance. Recently, Bhat *et al.* [29] proposed a novel formulation for predicting distance values by means of classification. Tian *et al.* [30] proposed attention blocks within the decoder, while other methods adopt a fully Transformer based architecture [31], [32].

b) Occlusion boundaries: All of the above works aim to learn a globally consistent depth map, yet, do not focus on fine local details, often resulting in blurred boundaries and deformed planar surfaces. Consistent with our work, Hu *et al.* [33] focuses on accurate occlusion boundaries through gradient and normal based losses. Ramamonjisoa *et al.* [25] aims to improve predicted depth boundaries by estimating normals and edges along with depth and establishing consensus between them. In a follow-up work, they apply wavelet decomposition at different scales and re-weight the feature map by a binary mask calculated at different frequency coefficient thresholds [34]. Several works apply bilateral filters to increase occlusion gaps [35], infer each object separately [36], or learn energy-based image-driven refinement focusing on edges [37], [16]. Cliffnet [38], combines multi-scale features to obtain sharper depth maps, whereas, Tosi *et al.* [39] tackles the issue through mixture density networks, focusing on stereo matching problem.

c) Sparse convolutions: Sparsity in convolutions has been investigated in several works [41], [42] aimed at improving the efficiency of neural networks by reducing the number of parameters, *i.e.* increasing sparsity. Minkowski Engine was proposed as an efficient 3D Spatio-Temporal convolution built on sparse convolutions [43]. In contrast to such works, Uhrig *et al.* used sparse meshes [17] to improve structural understanding in the case of sparse inputs, *e.g.*

a) Overall architecture



b) Instance Convolution

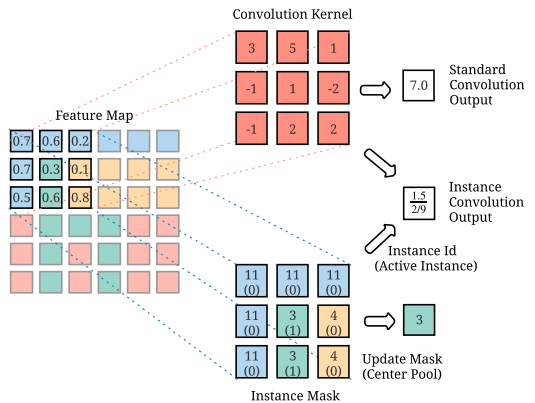


Fig. 3. **a) Schematic overview of proposed architecture.** The input image is fed into a state-of-the-art depth prediction model (e.g. SharpNet [25] or BTS [27]) to obtain global image features. The extracted depth features along with the object masks are then passed into the Instance Convolution block to predict a sharp depth map. **b) Masking mechanism of Instance Convolutions.** The abstract image on the left part contains a chair represented in green. The feature values of the masked region of the kernel are 0.3 and 0.6. The mask for the kernel shows the object pixels. If the normal convolution with a binary mask multiplication is used for this example, 0.9 can be obtained as the result, but with the Instance Convolution, the result is $\frac{0.9}{2/9}$. The successive mask is calculated by taking the center value of the current kernel mask.

depth map completion [44]. Some works used a similar convolution operator, called *partial* or *gated* convolution, for image editing and inpainting tasks [45], [46] to discard content-free regions. Su *et al.* [47] applies per-pixel kernel weight adaptation through a Gaussian assumption, omitting the boundaries. In our work, we are also interested in computing convolution only on a subset of input pixels. Differently from these works, our masks do not define a random set or a normally distributed sparse set of pixels. Our masks change dynamically for each pixel position, to extract features within the same *object segment*.

d) Over-segmentation methods: The proposed instance convolution operation relies on the detection of meaningful object segments in a scene. One alternative would be to identify all objects in a scene, either as annotated instance labels, or via learned segmentation models, e.g. Mask R-CNN [48]. The former requires a heavy amount of annotation for large datasets. The latter requires the objects in the corresponding dataset to match the pre-trained models and can additionally lead to inaccurate edges. To detect objects without pre-trained models and labeled data, in this work, we leverage over-segmentation methods. Among the available methods for over-segmentation [18], [40], [49], in our experiments we mainly focus on super-pixel (SLIC) [18] and Bayesian adaptive super-pixel segmentation (BASS) [40] (see Fig. 2).

III. METHODOLOGY

In this section, the problem statement and the individual components of the proposed method for boundary-aware MDP are presented.

A. Depth Estimation Using Deep Learning

Monocular depth estimation has recently received a lot of attention in literature and several methods have been proposed [11], [12], [10]. Interestingly, even very early methods have noticed the performance degradation around occlusion

boundaries, and various different measures, such as skip-connection [22], [23] or Conditional Random Fields [24], [50], have been put in place to counteract the smearing effect. Nevertheless, despite those measures, the proposed methods still can not capture the high frequencies of object discontinuities due to the inherent nature of 2D convolutions.

The classical convolution kernel simultaneously operates on all inputs within the kernel region, performing a weighted summation. Consequently, features originating from different object parts are simply fused, which in turn causes a blurring of the object boundaries, *i.e.* the edges that separate the object from the background in 3D. A corresponding example is shown in Fig. 1. Notice that this effect is more visible when viewing the associated 3D point cloud.

B. Instance Convolutions for Boundary-Aware Depth Estimation

To avoid aggregation of features appertaining to different image layers, we thus propose to leverage super-pixels in an effort to guide the convolution operator. In particular, inspired by Sparsity Invariant Convolutions [17], we propose *Instance Convolutions*, which applies the weighted summation only to pixels belonging to the same segment as the central pixel. In other words, it renormalizes the convolution function by considering the number of instance pixels that contributed to the output to keep it at a similar range than the common convolutional operator, as depicted in Fig. 4. Notice that when all pixels are within the instance mask, our formulation simply turns into standard convolution. Formally, for the n -th output feature map this can be written as

$$f_{u,v,n}(S, X, W) = \frac{\sum_i (S_{u,v} \odot X_{u,v} \odot W_n)_i}{\sum_{k_w * k_h} (S_{u,v})_i} + b, \quad (1)$$

with \odot denoting the element-wise multiplication, $X_{u,v}$ being the features within the window of the kernel of size $k_w \times k_h$

placed at (u, v) , $S_{u,v}$ being the indicator kernel which is 1 for each pixel within the window if it belongs to the same class as the central pixel (u, v) and 0 otherwise, and W_n denoting the applied kernel weights. All matrices $S_{u,v}, X_{u,v}, W_n \in \mathbb{N}^{d_{in} \times k_w \times k_h}$ with the indicator kernel $S_{u,v}$ being replicated, accordingly.

Since our architecture follows the standard encoder-decoder methodology (Section III-C), we have to adequately propagate the segment information through the network. However, as MaxPooling can lead to loss of spatial information, we introduce the center-pooling operator, which simply forwards the segment id of the central pixel with respect to each downsampling operation to preserve the object boundaries. Whereas at upsampling, the original semantic map (or downsampled from previous layers) is directly used as they are already readily available or computed. For a detailed explanation of Instance Convolutions, see Fig. 3 (b).

A deep learning-based approach would simply transfer the problem of clear boundaries towards the segmenter. Such methods can also never cope with all objects classes in the wild. Moreover, our method does not require any semantic information, but rather only needs to understand what pixels belong to the same discontinuity-free object part. Hence, we instead rely on optimization-based approaches, *i.e.* SLIC [18] and BASS [40], to obtain the needed super-pixels. Thereby, these works can provide not only object boundaries but also self-occlusions within objects (see Fig. 2). Noteworthy, while SLIC requires to define the number of output segments in advance, BASS can find the optimal number of segments by itself, however at a higher computational cost.

C. End-to-end Architecture

To summarize, we model our Instance Convolution such that the method is particularly suited for estimating depth at object boundaries. Nevertheless, as an output pixel has never observed any feature outside of the segment it resides on, it is impossible for the model to predict metric depth due to the scale-distance ambiguity (*i.e.* a large object far away can have the same projection onto the image plane as a small object close by). Therefore, we harness a state-of-the-art MDP backbone to extract global information about the scene. We then feed the extracted features together with the obtained super-pixels to our Instance Convolution-driven network to estimate the final edge-aware depth. Since the backbone as well as our Instance Convolution block are fully differentiable, we can train the whole model end-to-end. Proposed method can be plugged together with different depth predictors. In this paper, we use SharpNet [25], BTS [27], and VNL [13] for feature extraction, to show the generalizability of Instance Convolutions.

D. Training Loss

Our training objective is composed of three terms, *i.e.* the depth loss L_1 , the gradient loss L_{grad} , and the normal loss L_{normal} . Notably, these are standard choices also used by [33], [51]. The L_1 -term is the main cost function to guide the learning of per-pixel distance values by taking the absolute

difference between the predicted depth d and the ground truth depth d^{GT} with:

$$L_1(d, d^{GT}) = \frac{1}{N} \sum_{i=1}^N |d_i^{GT} - d_i|. \quad (2)$$

Furthermore, the depth gradient loss is given as

$$L_{grad}(d, d^{GT}) = \frac{1}{N} \sum_{i=1}^N |\nabla_h d_i - \nabla_h d_i^{GT}| + |\nabla_v d_i - \nabla_v d_i^{GT}|, \quad (3)$$

and penalizes missing fine details. The horizontal (∇_h) and vertical (∇_v) gradients are computed via the Sobel operator.

Finally, in an effort to learn finer-level surface details, such as high-frequency fluctuations on small structures, we utilize the angular distance between per-pixel normals according to

$$L_{normal}(n, n^{GT}) = \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{\langle n_i, n_i^{GT} \rangle}{\|n_i\| \cdot \|n_i^{GT}\|} \right). \quad (4)$$

Thereby, a surface normal n_i of a pixel can be computed by the vertical and horizontal gradients of the depth map as

$$n_i^g = [-\nabla_x(d_i), -\nabla_y(d_i), 1]^T. \quad (5)$$

The total training loss is the sum of the three losses as

$$L = L_1 + L_{grad} + L_{normal}. \quad (6)$$

IV. EXPERIMENTS

In this section, the experimental setup along with the proof-of-concept experiment is introduced.

A. Evaluation Criteria

Standard metrics. We follow the standard MDP metrics as introduced in [21] and report results as mean absolute relative error (*absrel*), root mean squared error (*rmse*), and the accuracy under threshold ($\delta_i < 1.25^i = 1, 2, 3$).

Occlusion boundaries. Since standard metrics evaluate depth distances in an averaged manner, Koch *et al.* [6] proposed another set of metrics to evaluate finer details in terms of occlusion boundaries and planarity. The Depth Boundary Error (DBE) calculates the accuracy (ϵ_{acc}) and completeness (ϵ_{com}) of occlusion boundaries by comparing predicted depth map edges to the annotated map of occlusion boundaries. The Truncated Chamfer Distance (TCD) is then calculated as:

$$\epsilon_{DBE}^{acc} = \frac{1}{\sum_i Y_i} \sum_i E_i \cdot Y_i, \quad (7)$$

where Y_i is the distance between the i 's predicted depth map's edge pixel, extracted by Canny edge detector, to the corresponding nearest ground truth edge pixel. If this distance is greater than 10 pixels, E_i amounts to zero to reject mismatched pixels. Oppositely, completeness ϵ_{DBE}^{com} is calculated by calculating the distances from ground truth edges to prediction to measure how much of the existing boundaries are measured. Furthermore, Directed Depth Error (DDE) assesses whether the overall prediction falls too short

or too far, by calculating the percentage of depth values that lie behind (ϵ^-), in front (ϵ^+), and in the proximity (ϵ^0) of a predefined median plane. Finally, the Planarity Error (PE) denotes the surface normal error on planar region maps computed as 3D shift distance (ϵ^{plan}) and angular difference (ϵ^{orie}). As we focus on occlusion boundary quality, we mainly consider DBE along with the standard metrics. We state the results for PE and DDE for completeness, and refer the reader to [6] for further details on their formulation.

B. Datasets

NYU v2. It consists of images collected in real indoor scenes [7]. Depth is captured with a Microsoft Kinect camera. The raw dataset of RGB depth pairs (approximately 120K images) has no semantic labels. The authors created a smaller split with semantic and instance labels, along with the refined depths and normals. In our experiments, we refer to this smaller split, which contains 1449 images in total, namely 795 for training and 654 for testing.

NYU v2 - OB. Occlusion boundary annotations on the NYU v2 test data for evaluation purposes is released by Ramamonjisoa *et al.* [16] within this dataset. We use it to evaluate the occlusion boundaries of our depth predictions.

iBims. This dataset [6] is presented as a set of around 100 images (for evaluation only) along with novel metrics on occlusion boundaries and planarity scores. They provide rich annotations of dense depth maps from different scenes, with occlusion boundaries and planar regions.

KITTI. Finally, this dataset provides a monocular depth prediction benchmark [19]. There are no occlusion boundary annotations given, so it is not possible to evaluate DBE scores. Hence, we only provide quantitative results *w.r.t* the standard metrics and employ qualitative samples to show the effectiveness of our method on outdoor scenes. In the following experiments, we trained models with 5000 sample images and took the original test split from KITTI-Eigen [19] for testing (629 images).

C. Overfitting Experiment

To test the effectiveness of our Instance Convolutions (IC), we first conduct overfitting experiments, comparing classical convolution based MDP methods against their IC counterpart. An exemplary outcome for this experiment is presented in Fig. 1. Here, the classical convolutions achieve an *absrel* of 0.009 and DBE ϵ_{acc} of 0.52, whereas our proposed IC reports 0.007 and 0.44 for *absrel* and DBE ϵ_{acc} . This confirms our expectation that IC has the capacity to prevent the boundary smoothing issue of standard convolutions, by only considering the relevant pixels within the segments.

D. Comparison to State-of-The-Art

In Table I, we compare our results from NYU v2 with three state-of-the-art approaches, namely SharpNet [25], VNL [13] and BTS [27]. Thereby, the proposed architecture (Fig. 3) uses these pre-trained models for latent depth feature extraction and applies Instance Convolution based blocks.

All models are trained using PyTorch on a NVidia RTX 2080Ti for 35 epochs with Adam optimizer, started learning

TABLE I
NYU v2 TEST SPLIT STATE-OF-THE-ART COMPARISON FOR WITH AND WITHOUT INSTANCE CONVOLUTION.

Method	Error ↓		Accuracy ↑			DBE ↓	
	<i>absrel</i>	<i>rmse</i>	δ_1	δ_2	δ_3	ϵ_{acc}	ϵ_{com}
SharpNet [25]	0.116	0.448	0.853	0.970	0.993	3.041	8.692
+ Instance Conv.	0.124	0.456	0.847	0.971	0.993	1.961	6.489
VNL [13]	0.112	0.417	0.880	0.975	0.994	1.854	7.188
+ Instance Conv.	0.117	0.425	0.863	0.970	0.991	1.780	6.059
BTS [27]	0.110	0.392	0.885	0.978	0.994	2.090	5.820
+ Instance Conv.	0.121	0.467	0.848	0.964	0.993	1.817	6.197

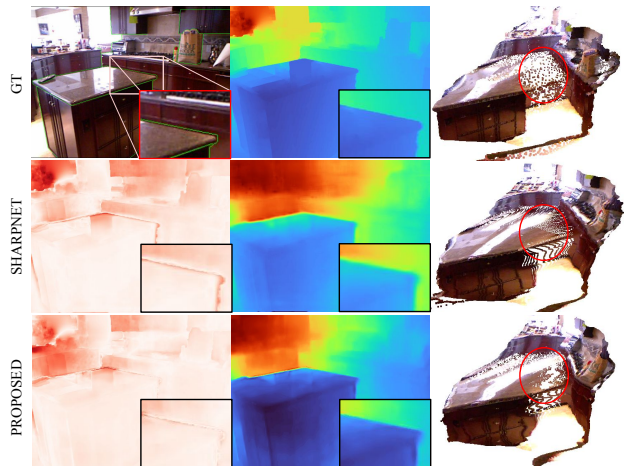


Fig. 4. **Error map comparison in NYUv2**, (top) the ground truth, (center) the SharpNet baseline [25], (bottom) the proposed IC model. The results are given as error maps, depth maps and point clouds. The improvement on the occlusion boundary is best visible when zoomed in.

rate with $1e - 4$ and decreased by 0.1 every 10 epochs. The loss terms in Eq. 6 have equal weights of 1. We use SLIC [18] to obtain super-pixels with 64 segments and set sigma to 1. SharpNet [25] is trained with a batch size of 4, extracting the feature embedding from UpConv1 layer having a shape of $128 \times 120 \times 160$. For BTS [27] we use a batch size of 3, and forward the output of the iConv1 feature layer of size $32 \times 480 \times 640$. Finally, VNL [13] also utilizes a batch size of 3, and propagates the feature from the final layer of shape $256 \times 480 \times 640$. Our proposed model then appends 3 layers of Instance Convolutions with gradually decreasing number of feature channels. The feature map resolution remains constant with a final kernel of size 1×1 .

We outperform the original methods with respect to the occlusion boundary metrics, and report comparable results for classical metrics. The qualitative results in the Fig. 5 further supports these findings, where the proposed Instance Convolution predictions of each model have sharper occlusion boundaries, and resulting depth maps. Further, in Fig. 4 we show an additional qualitative example to highlight the error distribution in a depth map. We observe that the error map of the SharpNet baseline (center row) has higher values along the occlusion boundaries, while ours (bottom row) are mostly uniform (*c.f.* zoomed-in region).

TABLE II
 QUANTITATIVE RESULTS ON IBIMS FOR STANDARD METRICS AND PE, DBE, AND DDE METRICS FOR DIFFERENT MDP METHODS.

Method	Error ↓			Accuracy ↑			PE (in $cm/^\circ$) ↓		DBE (in px) ↓		DDE (in %)		
	$absrel$	log_{10}	$rmse$	δ_1	δ_2	δ_3	ϵ^{plan}	ϵ^{orie}	ϵ^{acc}	ϵ^{comp}	$\epsilon^0 \uparrow$	$\epsilon^- \downarrow$	$\epsilon^+ \downarrow$
Eigen [21]	0.32	0.17	1.55	0.36	0.65	0.84	7.70	24.91	9.97	9.99	70.37	27.42	2.22
Laina [23]	0.26	0.13	1.20	0.50	0.78	0.91	6.46	19.13	6.19	9.17	81.02	17.01	1.97
Liu [50]	0.30	0.13	1.26	0.48	0.78	0.91	8.45	28.69	2.42	7.11	79.70	14.16	6.14
Li [52]	0.22	0.11	1.09	0.58	0.85	0.94	7.82	22.20	3.90	8.17	83.71	13.20	3.09
Liu [53]	0.29	0.17	1.45	0.41	0.70	0.86	7.26	17.24	4.84	8.86	71.24	28.36	0.40
SharpNet [25]	0.26	0.11	1.07	0.59	0.84	0.94	9.95	25.67	3.52	7.61	84.03	9.48	6.49
with Instance Conv.	0.29	0.12	1.14	0.55	0.82	0.92	9.83	25.88	3.11	7.83	81.84	8.27	9.88
BTS [27]	0.24	0.12	1.08	0.53	0.84	0.94	7.24	20.51	2.50	5.81	82.24	15.50	2.27
with Instance Conv.	0.22	0.11	1.11	0.57	0.86	0.94	6.76	19.39	3.71	8.01	84.04	13.3	2.67
VNL [13]	0.24	0.11	1.06	0.54	0.84	0.93	5.73	16.91	3.64	7.06	82.72	13.91	3.36
with Instance Conv.	0.23	0.10	1.06	0.58	0.85	0.93	5.62	16.53	3.03	7.68	83.85	13.26	2.87

In Table II, we further report our results with respect to the iBims evaluation dataset, in order to assess the generalizability of our method. Note that this dataset is used only for inference (*i.e.* no training), to measure whether the models are capable of detecting depth values. Here, in most cases our models again improved the counterpart depth models in terms of DBE accuracy and completion metrics. We additionally provide the qualitative results for iBims samples in Fig. 5 third column, with the highlighted occlusion boundary regions.

Notably, on iBims, the model with the best $absrel$ value does not have the best DBE score (Li *et al.* [52] 0.22 $absrel$, 3.90 DBE vs. Liu *et al.* [50] 0.30 $absrel$, 2.42 DBE). This can also be seen on NYU results, where our Instance Convolution improving DBE scores, also the occlusion boundaries as seen in qualitative illustrations, however being on-par on $absrel$ metrics. This conceptually agrees with the fact that $absrel$ averages out per pixel distances, while DBE calculates the drift errors on occlusion maps. This could explain the degradation on SharpNet $absrel$ score on iBims.

Finally, we evaluated our model on KITTI [19], extending BTS with a DenseNet-161 backbone [54]. Since KITTI does not provide annotations for occlusion boundaries, we mostly rely on qualitative evaluation to assess the performance in terms of boundary quality. Fig. 6 illustrates a few examples, comparing the BTS baseline with ours. Notably, our method not only produces sharper edges for the objects, but it is also able to recover sharp contours for the biker in the second image, which is completely missed by standard BTS, thanks to our over-segmentation formulation. Further, we are capable of producing sharper results without any loss in quality. In particular, using the standard metrics BTS reports an error of 0.062 and 2.798 for $absrel$ and $rmse$, respectively, while we instead report an error of 0.06 and 2.479.

E. Ablation Studies

In Table III, we present an ablation for different parameters and configurations using SharpNet [25] as backbone. Instance convolution based models are shown in the table

with IC or Standard Convolutions as SC, following by the number of segmented regions. The baseline model is the original SharpNet model without SC/IC.

Super-pixels information. We trained a model with SC and provided the super-pixel segmentation map as an extra input to each convolutional layer. As shown in Table III line SC 64, DBE score is better than the baseline, however, worse than the proposed IC. It shows that the proposed IC method is indeed able to use the super-pixels information as desired to reason on occlusion boundaries.

Number of segments in SLIC. We ablate different number of super-pixel segments. As expected, increasing number of segments improve DBE accuracy, yet, induce a little loss in $absrel$. Notice that this is also the case for most sota works. Best performing works for edges are often worse on $absrel$, which could be caused by imperfect annotations.

Over segmentation with BASS. We also qualitatively evaluated our methods using BASS [40] to extract super-pixels. As shown in the Fig. 2, BASS is able to retrieve more detailed segments from the image, however it also detects overly noisy edges due to redundant number of segments (200-250), which increases the model complexity and makes the learning more difficult.

Instance masks. To compare the quality of instance mask prediction to unsupervised segmentation, we ablated our method with the state-of-the-art instance mask prediction method PointRend [55]. As both the $absrel$ and the DBE results were poorer than the baseline, this proved the effectiveness of over-segmentation method, most likely because of detecting self-occlusions within the images. Additionally, we show our results using ground truth instance masks, which, as expected, improve performance over predicted masks. Nevertheless, very interestingly, our proposed IC approach with 128 segments performs a little better than ground truth masks *w.r.t.* ϵ_{acc} , which we contribute to the fact the instance masks neglect intra-shape boundaries (see Fig. 2).

Runtime analysis Full inference times are given in Table III as Frames per Second (FPS) excluding segmentation calculation, FPS* including. Each runtime is averaged on

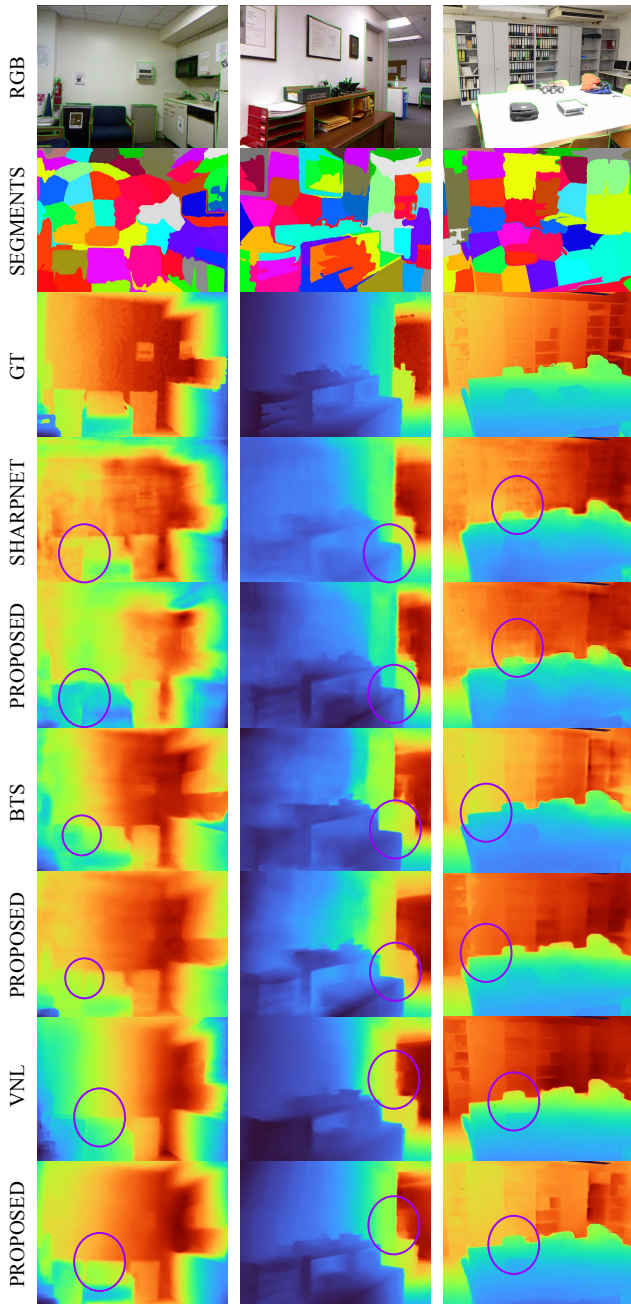


Fig. 5. **Qualitative results.** The first two columns are samples from NYUv2 and last column from iBims dataset, rows respectively show the RGB input image, SLIC output for the input image, ground truth depth, and depth maps from original SharpNet [25], BTS [27] and VNL [13] models, and the improved depth maps from proposed Instance Convolution models.

1000 inferences. IC does not excessively alter the FPS compared to SharpNet.

V. CONCLUSIONS

In this work, we introduce a novel depth estimation method, which is particularly tailored towards tackling the problem of depth smoothing at object-discontinuities. To this end, we propose a novel convolutional operator, which avoids feature aggregation across discontinuities by means of super-pixels. Our exhaustive evaluation on NYU depth v2, iBims

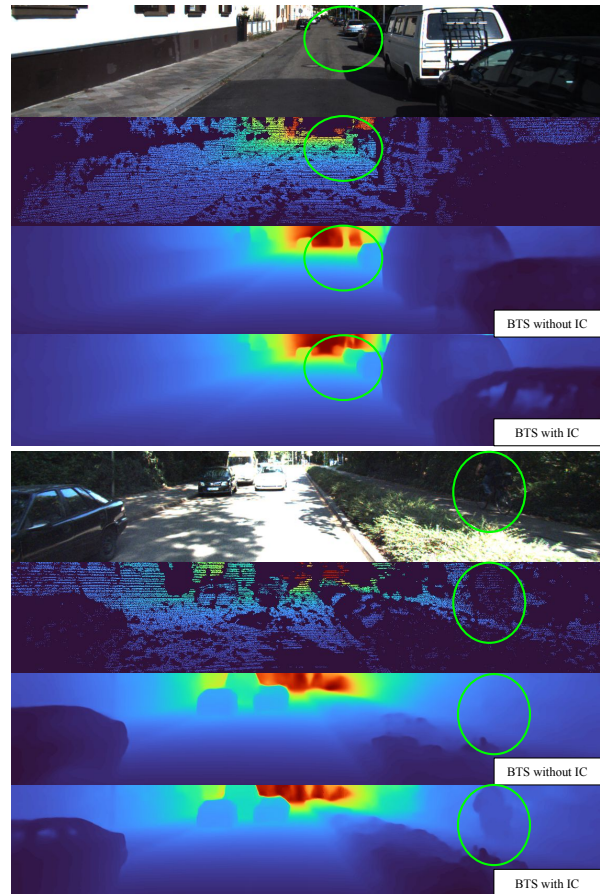


Fig. 6. **Qualitative results on KITTI [19]**, using the BTS [27] backbone, showing (from top to bottom) the input RGB, ground truth depth map, BTS prediction, and proposed IC model prediction. The improvements on occlusion boundaries around objects are highlighted with green-circles.

TABLE III

ABLATION STUDY ON NYUV2, COMPARING USAGE OF DIFFERENT MASKS (GROUND TRUTH, POINTREND, AND BASS), SUPER-PIXELS WITH STANDARD CONVOLUTIONS (SC), AND DIFFERENT NUMBER OF SEGMENTS (16-32-64-128) WITH INSTANCE CONVOLUTIONS (IC).

Method	Error ↓		DBE ↓		Runtime	
	<i>absrel</i>	<i>rmse</i>	ϵ_{acc}	ϵ_{com}	FPS	FPS*
SharpNet [25]	0.12	0.45	3.04	8.69	16.7	16.7
GT Masks	0.12	0.46	2.05	6.49	13.5	13.5
PointRend [55]	0.13	0.45	2.21	6.76	13.5	3.64
BASS [40]	0.12	0.46	2.19	6.63	13.2	0.59
IC 16	0.14	0.47	2.07	6.59	13.5	3.08
IC 32	0.14	0.47	2.09	6.66	13.6	3.04
IC 64	0.12	0.46	1.96	6.48	13.4	2.97
SC 64	0.12	0.45	2.18	6.63	15.2	3.05
IC 128	0.13	0.46	1.92	6.57	13.3	2.89

as well as KITTI demonstrates that the proposed method is capable of improving depth inference around edges, while almost completely maintaining the quality on the remaining regions. In the future, we want to explore how Instance Convolution can be incorporated into other domains such as semantic segmentation to similarly improve sharpness.

REFERENCES

- [1] C. Veiga Almagro, M. Di Castro, G. Lunghe, R. Marín Prades, P. J. Sanz Valero, M. F. Pérez, and A. Masi, "Monocular robust depth estimation vision system for robotic tasks interventions in metallic targets," *Sensors*, vol. 19, no. 14, 2019. [Online]. Available: <https://www.mdpi.com/1424-8220/19/14/3220>
- [2] K. Tateno, F. Tombari, I. Laina, and N. Navab, "Cnn-slam: Real-time dense monocular slam with learned depth prediction," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017.
- [3] M.-L. Shih, S.-Y. Su, J. Kopf, and J.-B. Huang, "3d photography using context-aware layered depth inpainting," in *CVPR*, 2020.
- [4] A. Saxena, S. H. Chung, and A. Y. Ng, "Learning depth from single monocular images," in *NeurIPS*, ser. NIPS'05, 2005.
- [5] D. Hoiem, A. A. Efros, and M. Hebert, "Automatic photo pop-up," *ACM Trans. Graph.*, vol. 24, no. 3, p. 577–584, 2005.
- [6] T. Koch, L. Liebel, F. Fraundorfer, and M. Körner, "Evaluation of cnn-based single-image depth estimation methods," in *ECCV Workshops*, L. Leal-Taixé and S. Roth, Eds., 2019.
- [7] P. K. Nathan Silberman, Derek Hoiem and R. Fergus, "Indoor segmentation and support inference from rgbd images," in *ECCV*, 2012.
- [8] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *CVPR*, 2017.
- [9] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *CVPR*, 2016.
- [10] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *TPAMI*, 2020.
- [11] L. Wang, J. Zhang, O. Wang, Z. Lin, and H. Lu, "Sdc-depth: Semantic divide-and-conquer network for monocular depth estimation," in *CVPR*, 2020.
- [12] S. M. H. Miangoleh, S. Dille, L. Mai, S. Paris, and Y. Aksoy, "Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging," in *CVPR*, 2021.
- [13] W. Yin, Y. Liu, C. Shen, and Y. Yan, "Enforcing geometric constraints of virtual normal for depth prediction," in *ICCV*, 2019.
- [14] Y. S. . R. L. Jianbo Jiao, Ying Cao, "Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss," in *ECCV*, 2018.
- [15] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao, "Deep Ordinal Regression Network for Monocular Depth Estimation," in *CVPR*, 2018.
- [16] M. Ramamonjisoa, Y. Du, and V. Lepetit, "Predicting sharp and accurate occlusion boundaries in monocular depth estimation using displacement fields," in *CVPR*, 2020.
- [17] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity invariant cnns," in *3DV*, 2017.
- [18] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *TPAMI*, vol. 34, no. 11, pp. 2274–2282, 2012.
- [19] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," *CVPR*, 2012.
- [20] A. Saxena, M. Sun, and A. Y. Ng, "Make3d: Learning 3d scene structure from a single still image," *TPAMI*, vol. 31, no. 5, 2009.
- [21] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *NeurIPS*, 2014.
- [22] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," in *ICCV*, 2015.
- [23] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab, "Deeper depth prediction with fully convolutional residual networks," in *3DV*, 2016.
- [24] D. Xu, E. Ricci, W. Ouyang, X. Wang, and N. Sebe, "Multi-scale continuous crfs as sequential deep networks for monocular depth estimation," in *CVPR*, 2017.
- [25] M. Ramamonjisoa and V. Lepetit, "Sharpnet: Fast and accurate recovery of occluding contours in monocular depth estimation," *ICCV Workshops*, 2019.
- [26] Z. Zhang, Z. Cui, C. Xu, Y. Yan, N. Sebe, and J. Yang, "Pattern-affinitive propagation across depth, surface normal and semantic segmentation," in *CVPR*, 2019.
- [27] J. H. Lee, M.-K. Han, D. W. Ko, and I. H. Suh, "From big to small: Multi-scale local planar guidance for monocular depth estimation," *arXiv preprint arXiv:1907.10326*, 2019.
- [28] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *CVPR*, 2018.
- [29] S. Bhat, I. Alhashim, and P. Wonka, "Adabins: Depth estimation using adaptive bins," *CVPR*, 2021.
- [30] T. Chen, S. An, Y. Zhang, C. Ma, H. Wang, X. Guo, , and W. Zheng, "Improving monocular depth estimation by leveraging structural awareness and complementary datasets," *ECCV*, 2020.
- [31] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," *ICCV*, 2021.
- [32] G. Yang, H. Tang, M. Ding, N. Sebe, and E. Ricci, "Transformer-based attention networks for continuous pixel-wise prediction," in *ICCV*, 2021.
- [33] J. Hu, M. Ozay, Y. Zhang, and T. Okatani, "Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries," in *WACV*, 2019.
- [34] M. Ramamonjisoa, M. Firman, J. Watson, V. Lepetit, and D. Turmukhambetov, "Single image depth prediction with wavelet decomposition," in *CVPR*, 2021.
- [35] M.-L. Shih, S.-Y. Su, J. Kopf, and J.-B. Huang, "3d photography using context-aware layered depth inpainting," in *CVPR*, 2020.
- [36] H. Dhano, N. Navab, and F. Tombari, "Object-driven multi-layer scene decomposition from a single image," in *ICCV*, 2019.
- [37] S. Niklaus, L. Mai, J. Yang, and F. Liu, "3d ken burns effect from a single image," *ACM ToG*, vol. 38, no. 6, pp. 184:1–184:15, 2019.
- [38] L. Wang, J. Zhang, Y. Wang, H. Lu, and X. Ruan, "Clifnet for monocular depth estimation with hierarchical embedding loss," in *ECCV*, 2020.
- [39] F. Tosi, Y. Liao, C. Schmitt, and A. Geiger, "Smd-nets: Stereo mixture density networks," *CVPR*, 2021.
- [40] R. Uziel, M. Ronen, and O. Freifeld, "Bayesian adaptive superpixel segmentation," in *ICCV*, 2019.
- [41] B. Liu, M. Wang, H. Foroosh, M. Tappen, and M. Pensky, "Sparse convolutional neural networks," in *CVPR*, 2015.
- [42] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," in *NeurIPS*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds., 2016.
- [43] C. Choy, J. Gwak, and S. Savarese, "4d spatio-temporal convnets: Minkowski convolutional neural networks," in *CVPR*, 2019.
- [44] S. Zhao, M. Gong, H. Fu, and D. Tao, "Adaptive context-aware multi-modal network for depth completion," *IEEE Transactions on Image Processing*, 2021.
- [45] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, "Image inpainting for irregular holes using partial convolutions," in *ECCV*, 2018.
- [46] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. Huang, "Free-form image inpainting with gated convolution," in *ICCV*, 2019.
- [47] H. Su, V. Jampani, D. Sun, O. Gallo, E. Learned-Miller, and J. Kautz, "Pixel-adaptive convolutional neural networks," in *CVPR*, 2019.
- [48] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *ICCV*, 2017.
- [49] A. Levinshstein, A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi, "Turbopixels: Fast superpixels using geometric flows," *IEEE TPAMI*, vol. 31, no. 12, pp. 2290–2297, 2009.
- [50] F. Liu, C. Shen, and G. Lin, "Deep convolutional neural fields for depth estimation from a single image," in *CVPR*, 2015.
- [51] X. Chen, X. Chen, and Z.-J. Zha, "Structure-aware residual pyramid network for monocular depth estimation," in *IJCAI*, 2019.
- [52] J. Li, R. Klein, and A. Yao, "A two-streamed network for estimating fine-scaled depth maps from single rgb images," in *ICCV*, 2017.
- [53] C. Liu, J. Yang, D. Ceylan, E. Yumer, and Y. Furukawa, "Planetnet: Piece-wise planar reconstruction from a single rgb image," in *CVPR*, 2018.
- [54] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *CVPR*, 2017.
- [55] A. Kirillov, Y. Wu, K. He, and R. Girshick, "Pointrend: Image segmentation as rendering," in *CVPR*, 2020.