

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Semantics-Depth-Symbiosis: Deeply Coupled Semi-Supervised Learning of Semantics and Depth

Nitin Bansal¹, Pan Ji^{*1}, Junsong Yuan², Yi Xu¹ ¹OPPO US Research Center, USA ²State University of New York at Buffalo, USA

{nitin.bansal,yi.xu}@innopeaktech.com, peterji1990@gmail.com, jsyuan@buffalo.edu

Abstract

Multi-task learning (MTL) paradigm focuses on jointly learning two or more tasks, aiming for an improvement w.r.t model's generalizability, performance, and training/inference memory footprint. The aforementioned benefits become ever so indispensable in the case of training for vision-related dense prediction tasks. In this work, we tackle the MTL problem of two dense tasks, i.e., semantic segmentation and depth estimation, and present a novel attention module called Cross-Channel Attention Module (CCAM), which facilitates effective feature sharing along each channel between the two tasks, leading to mutual performance gain with a negligible increase in trainable parameters. In a symbiotic spirit, we also formulate novel data augmentations for the semantic segmentation task using predicted depth called AffineMix, and one using predicted semantics called ColorAug, for depth estimation task. Finally, we validate the performance gain of the proposed method on the Cityscapes and ScanNet dataset. which helps us achieve state-of-the-art results for a semisupervised joint model based on depth estimation and semantic segmentation.

1. Introduction

Convolutional Neural Networks (CNNs) [35] have helped achieve state-of-the-art results for a range of computer vision tasks including image classification [23], semantic segmentation [46, 3, 4, 37], and depth estimation [20]. Generally, each of these tasks is trained in isolation, assuming that inter-task features are largely independent. On the contrary, multiple works in the domain of multi-task learning (MTL) [67, 43, 80, 7, 29, 32, 57, 25] point towards an inherent symbiotic relation among multiple tasks, where one task benefits from other *sibling* tasks. In [66], Standley *et al.* particularly point towards a synergy between semantic segmentation and depth estimation task, when trained jointly.

Despite previous success, the lack of sufficient labeled data for multi-task training can affect the performance of MTL. To address this limitation, MTL models such as [24, 33, 47, 16] propose parameter sharing to overcome data sparsity and enforce generalization by leveraging task losses to regularize each other. The data sparsity problem is more emphatically seen for dense tasks such as semantic segmentation and depth estimation, where obtaining perfect per-pixel annotations is both expensive and untenable in some scenarios, thus fully supervised learning may not be feasible. Motivated by the above observations, we propose to leverage large-scale video data using a semi-supervised MTL training paradigm [34, 68, 52, 76] where semantic segmentation follows a semi-supervised setting and the depth sub-model is trained in self-supervised manner. Specifically, to improve semi-supervised semantic training, we propose an AffineMix data augmentation strategy, which aims to create new labeled images under a varied range of depth scales. Under this scheme, randomly selected movable objects are projected over the same image, for a randomly selected depth scale. This unlocks another degree of freedom in data augmentation scheme, generating images which are not only close to original data distribution but also more diverse and class balanced. On the depth augmentation end, we propose a simple yet effective data augmentation scheme called ColorAug, which establishes strong contrast between movable objects and adjacent regions, using intermediate semantics information. At last, we employ orthogonal regularization as an strategy to improve MTL training efficacy, which has a positive impact on both semantics and depth evaluation.

Besides training paradigm, the architecture of MTL model is equally important as it determines how exactly the intermediate features, belonging to different tasks, interact with each other. As highlighted in [80], *what to share and how to share* still remains an open question. To improve MTL performance, we adopt a hybrid parameter

^{*}Worked when associated with OPPO US Research Center

sharing approach [77, 25], which enforce a soft parameter sharing at decoder layers to facilitate both flexibility and inter-feature learnability. To emphasize inter-channel interactions between tasks, our proposed Cross-Channel Attention Module (CCAM) enforces dual attention on intermediate depth and segmentation features over both spatial and channel dimensions. This enables us to estimate a degreeof-affinity between inter task channel-features as an intermediary score in an end-to-end differentiable framework. Using CCAM, we can linearly weight the contribution of features from each task before sharing and thus encourage a more informed and reliable feature transfer between two tasks.

To summarize our main contributions:

- To improve feature sharing for MTL, we propose CCAM to estimate cross channel affinity scores between task features. Which enables better inter-task feature transfer.
- To deal with data sparsity for MTL, we design a dual data augmentation mechanism for both semantics and depth. Our method encourages diversity, class balance for semantic segmentation, and better region discrimination for depth estimation.
- We incorporate orthogonal regularization for depth and semantics with diminishing weighting to facilitate better feature generalization and independence.

2. Related Work

2.1. Multi-Task Learning (MTL)

MTL [2] has been used across various tasks, as it improves generalization by transferring domain information of related tasks as an inductive bias. Previous works such as [2, 24, 33, 42] have advocated for hard parameter sharing, where as others [47, 16, 49] advocate soft sharing of parameters, where specific parameters are connected through a learnable link module. For encoder-decoder architectures, previous works such as [60, 49, 16] are encoder focused, while others [81, 77, 25] mainly focus on the decoder modules. Our work follows a hybrid approach. We use a shared encoder but employ soft parameter sharing over the decoder module as shown in Fig 1(a). Our approach gravitates more towards the decoder, but improves from the encoderfocused approaches by enabling an efficient cross-task feature sharing mechanism using CCAM module. Work such as [44, 58] employ Residual adapter module, to dedicate a set of model's parameters to each task in a MTL setup; while we improve model's shared representation capability through effective inter-task feature sharing. From optimization point of view, works such as [8, 62, 22] enforce task balancing by using gradient based multi-objective optimization. We find little to no effect employing such regularization during our training. Finally both adversarial training [65, 48] and uncertainty [30] are utilized for MTL. Both of them have entirely different motivations compared to our work, and are not necessarily looking to better feature sharing, for two related tasks.

A few other methods [7, 57] leverage semantic labels to jointly train on depth and semantics but mainly to improve on depth results. Marwin et al. [32] focus on improving depth results specifically for moving objects in a scene using the intermediary semantics information, whereas Jiao et al. [29] tackle long-tail distribution in depth prediction domain using semantics. Some of these models are not truly self-supervised for depth estimation or fail to jointly improve both tasks. Another approach uses knowledge distillation [21, 77] from segmentation to guide depth estimation. Furthermore, some other methods [77, 43] exploit spatial attention [53, 71] for more effective cross-feature distillation and task related feature extraction. Most recently, Hoyer et al. [25] propose a joint training network, which has a feature sharing network between the task as proposed in [77]. Our experiments show that there is little impact on the result of depth metrics upon inclusion/exclusion of this feature sharing module, suggesting ineffective feature transfer between the two tasks. Also [25] focuses only on improving semantic results using depth estimation procedure but not viceversa. In contrast, CCAM facilitates proportional sharing of features between depth and semantics across each channel (Fig. 2). Please refer to Sec.3.2 for more details.

2.2. Semi-supervised Semantic Segmentation

Deep convolutional neural network models [46, 63] have been the go-to network for both supervised and semisupervised semantic segmentation tasks. Subsequent models have improved by leveraging multi-scale input images [11, 5, 13, 40, 41], which capture finer details using multi-scale features. Furthermore, there have been models using feature pyramid spatial pooling [45, 82] and altrous convolutions [3, 4, 6, 37, 78] to further assist in better perpixel feature learning and achieve state-of-the-art results. We choose an architecture similar to U-Net [59], details about which is in Sec. 4.

Semi-supervised semantic segmentation training makes use of an unlabeled set of data. Many approaches take image-level labels [36, 38, 74] and class activation maps [83, 72] as a weak supervision signal, which gives marginal assistance in a dense prediction task such as perpixel segmentation. Methods based on consistency training [34, 68, 52, 76] use the idea that label space for unlabeled data should remain broadly unchanged after adding noise or perturbation to the input. Ouali *et al.* [55] use the same idea but apply perturbation on encoder features instead. Similarly, CutMix [79] vouches for stronger augmentations, where crops from input images and pseudo la-



Figure 1. (a): Block diagram of the complete multi-task model with semantics and depth losses. (b): Orthogonal regularization enforced on shared encoder(SE), depth decoder(DD) and semantic decoder(SD).

bels are used to generate additional training data. Class-Mix [54] takes it a step forward by using pseudo labels to get a mix mask, which is then used in consistency training. Our proposed data augmentation is most similar to DepthMix [25], which uses the idea of ClassMix [54] but also maintains geometric consistency. Our approach differs from DepthMix [25] on three counts. Firstly, we propose a new data augmentation strategy which generates the pseudo labels for selected foreground classes under different randomly selected depth scales, keeping geometric consistency intact. Secondly, we mix foreground masks over the same image and not across the other images for a given batch. Lastly we only consider movable objects as part of the data augmentation with the idea of better handling the intrinsic data bias due to class imbalance. Please refer to Sec. 3.3 for more details.

2.3. Self-supervised Depth Estimation (SDE)

Depth estimation in the absence of per-pixel ground truth is a well-studied problem. The self-supervised model relies on minimizing the image reconstruction loss, for an input that could either be in stereo-pairs or in a monocular sequence format. Depth estimation under a stereo setting [20, 17, 19] mainly focuses on predicting pixel disparity between the pairs and enforcing a consistency between left and right views. In the monocular sequence scenario, methods such as [20, 84, 17, 19, 27, 85, 69, 28] minimize the photometric reprojection loss during the training phase using the predicted depth and pose. Our depth module largely follows Godard *et al.* [20]. In addition to the reprojection loss, we also enforce a per-pixel minimum appearance loss and auto-masking which further improves prediction for occluded and stationary pixels.

3. Proposed Method

In this section, we start with the basic architecture in Sec. 3.1. We then discuss in detail about building blocks of CCAM in Sec. 3.2. Different strategies used for data augmentation for semi-supervised semantics and self-supervised depth network are subsequently discussed in Sec. 3.3. We then briefly discuss the training strategy for depth and semantics part in Sec. 3.4 and Sec. 3.5 respectively. We conclude this section by presenting how we can effectively apply orthogonal regularization on semantics and depth modules together to further enhance model's performance in Sec. 3.6.

3.1. Basic Architecture

For the multi-task training at hand, we follow soft or partial parameter sharing between the tasks. We use a shared encoder but separate decoders for semantics and depth tasks respectively as shown in Fig. 1. Apart from this, we use a separate encoder block for camera pose estimation, and an encoder pretrained on ImageNet [61] to calculate the feature distance loss, similar to [25]. The CCAM block (details in suppl. material) consists of mainly two sub-blocks, which mainly have convolutional, global average pooling and fully connected layers to compute spatial and cross-channel attention respectively. More specific details about different blocks are mentioned in Sec. 4.1, under network architecture.

3.2. Cross Attention Network Architecture

Cross-task feature transfer, can be broadly divided in to three sub-categories as described in [80]: (i) sharing initial layers to facilitate learning common features for complimentary tasks; (ii) using adversarial networks to learn common feature representation as in [64]; and (iii) learn-



Figure 2. Complete sub-blocks of Cross Channel Affinity Module (CCAM). N: Batch Size, C: Number of channels, H: Height, and W: Width.

ing different but related feature representations as presented in [50]. In this very context, Xu *et al.* [77] propose a multimodal distillation block, which shares cross-task features through message passing. It simulates a gating mechanism as shown in Eq. (1) and (2), by leveraging the spatial attention maps of each individual features of all tasks, which then helps decide what features a given task would share with other tasks. Without loss of generality, suppose we train a total number of T tasks, and $F_i^{\ k}$ denotes the i^{th} feature of the k^{th} task before message passing and $F_i^{\ o,k}$ after message passing. Xu *et al.* [77] define the message transfer as:

$$F_{i}^{o,k} = F_{i}^{k} + \sum_{t=1(\neq k)}^{T} G_{i}^{k} \odot (W_{t,k} \otimes F_{i}^{t}), \quad (1)$$

where \odot means element-wise product, \otimes represents convolution operation, $W_{t,k}$ represents the convolution block and $G_i^{\ k}$ denotes the gating matrix for i^{th} feature of the k^{th} task:

$$G_i^{\ k} = \sigma(W_g^{\ k} \otimes F_i^{\ k}) , \qquad (2)$$

where $W_a^{\ k}$ is a convolution block and σ denotes the sigmoid operator. We refer reader to [77] for more details regarding above message passing strategy. According to Eq. (1), it only shares cross-task features naively across the channel dimension. Suppose we are training simultaneously for two tasks namely: F^k and F^l . Eq. (1) indirectly implies that i^{th} channel-feature of F^k is only important to i^{th} channel-feature of F^l , which is not necessarily true in all scenarios. We overcome this major limitation by designing a module that calculates an affinity vector α_i , which gives an estimate about how the i^{th} channel of task F^k is related to any j^{th} channel of task F^l . As shown in Fig. 2, the entire process of building scores of inter-task channels can be subdivided into four sub-blocks. Since in our case we are dealing with two tasks, we denote them as Task A and Task B. We extract intermediate output features from respective decoder modules represented by A_F and B_F respectively, where $A_F, B_F \in \mathbb{R}^{N \times C \times H \times W}$.

CCAM Sub-block 1: We start with sub-block 1, where

task's intermediate features A_F and B_F are passed through a sequence of conv-blocks (W_A and W_B), which serves as spatial attention layers, to compute A_{SF} and B_{SF} according to the following equations:

$$A_{SF_i} = \sigma(W_A \otimes A_{F_i}) , \qquad (3)$$

$$B_{SF_i} = \sigma(W_B \otimes B_{F_i}) . \tag{4}$$

The idea is to get much more refined features from both tasks before estimating their cross-correlation. The output of this layer preserves the spatial resolution of the input features and gives output features represented by A_{SF} and B_{SF} respectively.

CCAM Sub-block 2: Subsequently in sub-block 2, we build a cross-task relation matrix C_{Mi} for each channel *i* of A_{SF} , where $C_{Mi} \in \mathbb{R}^{B \times C \times H \times W}$. We then pass the resultant matrix C_{Mi} to a channel attention module Ψ , which estimates the affinity vector α_i between i^{th} channel of A_{SF} and all the channels of B_{SF} in accordance with the equation:

$$\alpha_i = \Psi(A_{SF_i} \otimes (B_{SF})^T) , \qquad (5)$$

where Ψ denotes a combination of global average pooling layer followed by fully connected layers, with a sigmoid layer at the end, which serves as channel attention layer. We repeat this for all the channels of A_{SF} to get the corresponding affinity vector α .

CCAM Sub-block 3: As part of sub-block 3, we accumulate *Affinity Scores* for all the channels of A_F to achieve a final cross task affinity matrix M, given by the equation:

$$M = [\alpha_i \oplus \alpha_j] \quad \forall i, j \in [0, C) , \tag{6}$$

where \oplus denotes concatenation across the row dimension. **CCAM Sub-block 4:** Finally in sub-block 4, the cross task affinity matrix M achieved serves as a score accumulator, which helps get a linearly weighted features A'_F and B'_F , given by equation:

$$A'_{F} = A_{F} + (B_{F} \odot M), B'_{F} = B_{F} + (A_{F} \odot M^{T}),$$
(7)

where \odot represents element-wise multiplication.

3.3. Data Augmentation

Data augmentation plays a pivotal role in all machine learning tasks, as it helps gather varied data samples from a similar distribution. In the spirit of cooperative multitasking, in this subsection, we introduce novel ways of data augmentation on both segmentation and depth estimation tasks using predicted depth and semantics respectively.

Data Augmentation for Segmentation In the context of semi-supervised semantic segmentation, models such as [54, 79, 15] leverage consistency training by mixing image masks across two different images to generate a new image and its semantic labels. Hover et al. [25] go a step further to generate a much more diverse mixed label space by maintaining the integrity of the scene structure. We propose a new version of data augmentation called AffineMix, which considers mixing labels within the same image under a varied range of random depth values, thus producing a new set of affine-transformed images (see Fig 3(i)). To further improve the data augmentation process, we mix masks associated with only movable objects, to counter class imbalance, which is stark in a dataset such as Cityscapes. Given an image I and corresponding predicted depth D, we seek to generate a mixed image I', by scaling the depth of a selected movable object *m* by a scale factor of *s*:

$$D' = s * D , \qquad (8)$$

such that its spatial location in the image is changed in a geometrically realistic way. Changing the depth by a factor of *s*, results in an inverse scaling in the image domain and translational shift which would be given by:

$$t_x = (1.0 - 1/s) * o_x , \qquad (9)$$

$$t_y = (1.0 - 1/s) * o_y , \qquad (10)$$

where o_x and o_y are normalized offsets along x and y directions. Using t_x , t_y and inverse scaling 1/s, we can perform affine transformation on the image and label space to generate I_a and L_a . We then estimate the foreground mask, by comparing the new and old depths and masking it with the region which has the movable object in I_a and name it M_m . The final image and label would be then given by:

$$I' = M_m \odot I_a + (1 - M_m) \odot I , \qquad (11)$$

$$L' = M_m \odot L_a + (1 - M_m) \odot L . \tag{12}$$

Data Augmentation for Depth As shown in work [18, 75], factors such as position in the image, texture density, shading, and illumination are some of the pictorial cues about distance in a given image. Recent work in this field [12] also re-emphasizes the importance of contrast between adjacent regions as well as bright and dark regions within an

image. We particularly leverage this simple albeit important observation to develop an effective data augmentation technique called ColorAug, which uses different appearance based augmentation on movable and non-movable objects. We use the intermediate semantic labels predicted by the model to guide us in developing such an data augmentation as shown in Fig. 3(ii).

3.4. Self-Supervised Depth

Training formulation of the self-supervised depth network mainly follows good practices of [20] in terms of using a per-pixel minimum appearance, reprojection loss, and an auto-masking strategy. Main backbone of the depth network comprises of an encoder-decoder structure, represented by SE and DD in Fig. 1. We use a subsidiary pose network PN, to predict the relative translation (T) and rotation (R) of the source frames I_{t-1} and I_{t+1} with respect to the target frame I_t . Predicted poses and depth are then used to estimate the self-supervised depth loss, denoted by L_D . Building blocks of the encoder and decoder are specified in more details in Sec. 4.1.

3.5. Semi-Supervised Semantic Segmentation

For the semi-supervised semantic segmentation module, we start with a set of labeled images Ω_L , unlabeled images Ω_U , and N unlabeled image sequences. We pretrain the pose network PN, shared encoder SE, and depth decoder DD (see Fig. 1) modules with N unlabeled image sequences, in a similar fashion as mentioned in [25]. During the joint-training step, parameters of the depth decoder (DD) are used to initialize the segmentation decoder (SD)module. The CCAM module springs into action during this stage of training, facilitating informative features transfer between DD and SD modules. During semi-supervised training, suppose the labeled and unlabeled samples are represented as (I_L, S_L) and (I_U, S_U) , where S_U represents pseudo labels that are generated using a mean teacher algorithm [68]. The parameters of the mean teacher θ_T is given by the following equation:

$$\theta_T' = \alpha \theta_T + (1 - \alpha) \theta_{SD} , \qquad (13)$$

where θ_{SD} and α represent parameters of the segmentation decoder module and smoothing coefficient hyper-parameter respectively. We then generate the pseudo labels S_U as suggested in [54] as:

$$S_U = \arg\max_{I}(\theta_T(I_U)) , \qquad (14)$$

where l represents all possible classes in the Cityscapes dataset. Using Eq. (13) and (14), we can formalize the total semi-supervised loss as:

$$L_{SSL} = L_{CE}(\theta_{SD}(I_L), S_L) + L_{CE}(\theta_{SD}(I_U), S_U) .$$
(15)



Figure 3. (i) **AffineMix for Semantic Segmentation:** Steps involved : (a) GT/ Generated Psuedo label image; (b) Input image; (c) Randomly selected movable object (**'Person'** class) with random depth scaling of **0.75**; (d) Occlusion aware, affine generated foreground mask; (e) Augmented label image; (f) Augmented Image. (ii) **ColorAug for Depth:** Using intermediate semantics output, purposefully create regions of different brightness, contrast and saturation around movable objects highlighted using blue boxes. Row 1: Original Image, Row 2: Augmented Image.

We then use our AffineMix samples, as a substitute for I_U and S_U in Eq. (15) to get the final semi-supervised loss.

3.6. Orthogonal Regularization (OR)

Orthogonality constraint on model's parameters has shown encouraging results for tasks such as image classification [1, 70, 39], image retrieval [70], 3D classification [56] to name a few. Enforcing orthogonality has also helped improve model's convergence, training stability, and promote learning independent parameters. In a multi-task setup such as ours, feature independence within a given task is also important. We study the effect of applying a variation of the orthogonal scheme proposed in [1] on different submodules. The new loss function of the model, after adding orthogonal constraint is given by:

$$L_I = L_{SSL} + L_D, \tag{16}$$

$$L_F = L_I + \lambda \cdot \| W^T W - I \|_{\sigma} , \qquad (17)$$

where W, $\|\cdot\|_{\sigma}$, I, L_F and L_I represent the weights (for each layer), spectral norm, identity matrix, final model loss, and initial loss respectively. We find that enforcing orthogonality, particularly on the parameters of shared encoder *SE*, depth decoder *DD*, and segmentation decoder *SD* has the most positive impact on the model's performance. We confirm this by calculating average inter-channel correlation, for all decoder layers for both the tasks, with and without OR (more details in supplementary material). We postulate that independent features within semantics and depth module would make feature transfer between the tasks more effective. Details about enforcing this regularization and ablation study based on this are provided in Sec. 4.1.

Model	1/8(372)	1/4(744)	Full(2975)	
Adversarial[26]	58.80	62.30	_	
s4GAN [51]	59.30	61.90	65.80	
DST-CBC [14]	60.50	64.40	-	
CutMix[15]	60.34	63.87	67.68	
ClassMix [54]	61.25	63.30	-	
3-Ways [25]	68.01	69.38	71.16	
Ours.	70.72	71.65	72.93	

Table 1. Comparative mIoU (in %) numbers on Cityscapes Validation set. Reported numbers are mean over three runs, under same experimental settings.

4. Experiments

4.1. Experimental Details

Dataset We use the Cityscapes dataset [9] and ScanNet dataset [10] for training and evaluating the model. For evaluating on semantics we use the ground-truth labels provided as part of the datasets. We use the data prepossessing and data augmentation step as suggested in [25], where we downsample and center-crop the original 1024x512 color images to 512x512 images for training. For depth we use the unlabeled frames provided by the Cityscapes dataset during the training phase, whereas depth is evaluated against 1,525 images from 6 cities taken from the test set, for which we use the ground-truth depths generated by Watson *et al.* in [73], using the disparity maps and for ScanNet we continue to use the provided ground-truth. More details regarding it is provided in the suppl. material.

Network Architecture Basic network as part of our multitask training is similar to [25, 77] as seen in Fig. 1(a), which provides an intuitive and effective network for simultaneous training. It comprises of a shared encoder network which is ResNet101 [23] and two separate but identical architecture based on Deeplabv3 [3] with a U-Net [59] decoder. For pose network we use ResNet18 [23], whereas for the shared encoder network we use ResNet101 [23] pretrained on the ImageNet dataset [61]. We refer the readers to the original



Figure 4. **Qualitative results:** In each row we compare the semantics and depth results with the baseline [25]. Red boxes identify shortcomings in the baseline method whereas green boxes highlight the corresponding improvement by using our method.

Model	Seg. Metrics	Depth Metrics					
	mIoU↑	AbsRel↓	SqRel↓	RMSE↓	a1↑	a2↑	a3↑
3-ways* [25]	68.09	0.150	2.032	7.492	0.824	0.953	0.985
3-ways*[25] w/o attn	68.07	0.152	2.497	7.621	0.824	0.952	0.982
Ours (CCAM)	69.35	0.142	1.653	7.230	0.824	0.957	0.988
Ours + AM	69.84	0.149	1.651	7.521	0.817	0.952	0.984
Ours + AM + OR	70.72	0.146	1.546	7.239	0.815	0.953	0.987
Ours + AM + OR + CA	70.20	0.142	1.553	7.284	0.824	0.956	0.987

Table 2. Comparative mIoU and depth results between baseline model with (vanilla) attention, without attention, and with cross channel affinity attention. Models with * show the results as reproduced by us running the original model. AM: AffineMix, OR: orthogonal Regularization, CA: ColorAug.

Metric	CutOut	CutMix	ClassMix	DepthMix	AffineMix
mIoU	57.74	60.34	63.86	68.09	68.70

Table 3. Table presents the comparative performance of AffineMix method with previous semi-supervised works, which establishes Affine-Mix superiority over previous approaches on Cityscapes dataset(1/8 labeled images).

paper's [25] suppl. section for more details. CCAM block consists of mainly two sub-modules, which are, spatial and channel attention layers. Spatial attention comprises of convolutional blocks with kernel size of 3x3. Where as channel attention consist of a global average layer, and two fully connected layers respectively, followed by a sigmoid activation layer. Finer details about CCAM block architecture are provided in suppl. material.

Training For most part of the training, we follow the strategy as mentioned by Hoyer *et al.* [25]. We first train the self-supervised depth and pose module alone using Adam [31] till 300K iterations, with an initial learning rate of 10^{-4} , which is reduced by a factor of 10 using the step-scheduled learning rate mechanism. In the second iteration

we only look to finetune the shared encoder with a ImageNet feature distance for another 50K iterations. During joint training, we use SGD with an initial learning rate of 10^{-3} and 10^{-2} for encoder and decoder respectively, which are reduced by a factor of 10 after 30K iterations. We refer readers to [25] for more details. In our observation, we find freezing depth decoder parameters during data augmentation steps of self-supervised semantic learning leads to better model stability without any adverse effect on semantics or depth results. For orthogonal regularization we follow SRIP *et al.* [1] based method, starting with an initial weight $\lambda_O = 10^{-4}$, which is then reduced gradually to 10^{-5} , 10^{-6} , and 10^{-7} after 10K, 20K and 30K iterations respectively.

Model	Seg. Metrics	Depth Metrics					
	mIoU↑	AbsRel↓	SqRel↓	RMSE↓	a1↑	a2↑	a3↑
3-ways* [25]	39.05	0.184	0.106	7.492	0.688	0.931	0.986
Ours (CCAM)	41.57	0.174	0.095	7.230	0.715	0.939	0.989

Table 4. Comparative mIoU and depth results between baseline model with (vanilla) attention, and with cross channel affinity attention for ScanNet dataset

4.2. Ablation Studies

Cross Channel Affinity Block We study the effectiveness of the self-attention distillation module used in baseline [25] for cross-feature transfer between depth and semantics tasks. We follow that up with an experiment, incorporating the CCAM module and validate its positive impact on both tasks. As shown in Tab. 2, we observe minimal improvement to semantics and depth metrics with and without self-attention based distillation module; with mean IoU and absolute relative errors hovering around 68% and 15%. Including just the CCAM module, we see an improvement of **1.28%** in mIoU and a drop of **5.3%** in the absolute relative depth error.

AffineMix Augmentation We further incorporate the proposed AffineMix data augmentation, mainly with an aim to improve on semi-supervised semantic segmentation training. Through our experiments, we find that AffineMix consistently improves upon the baseline by **1.75%** respectively as shown in Tab. 2. We see a drop in the depth metrics after applying AffineMix. We postulate this drop in depth performance is due to visual incoherence in the local region adjoining to newly added objects, which might be detrimental to the depth task. We also did a comparative ablation experiment verifying the efficacy of proposed data augmentation with the previous data augmentation approaches. As shown in Tab. 3, we observe an improvement of about 8.36%, 4.64% and, 0.61% in mIoU numbers when compared to [79, 54, 25] respectively.

Orthogonal Regularization For verifying the efficacy of OR (Fig. 1(b)), we start out by comparing models trained with and without OR, from the perspective of feature independence after the completion of training. As postulated, we find that the features across all 4 layers of depth and semantics decoder module are much more independent after applying such a regularization.(More details in supplementary material). We further validate this claim by seeing a consistent improvement in semantics and depth metrics. We see an average gain of about **0.88%** over a run of three experiments for AffineMix augmentation, and an improvement of about **2.63%** compared to the baseline model. We also observe a marginal boost to depth metrics after inclusion of OR by **2.01%** (see Tab. 2).

ColorAug Augmentation For the depth augmentation experiment, we leverage the intermediate semantics information to create a new set of images, which help us create regions of different contrast, brightness, and saturation

around movable objects. We use this data augmentation scheme only after 10K steps of training, such that semantic outputs from the network are reliable to some extent (mIoU $\geq 60\%$). We achieve a reduction of about 2.8% to reach an absolute relative error of **14.2%** as seen in Tab. 2, along with an improvement of about **2.11%** for semantics.

We also conduct ablation experiments covering all different permutations of AM, CA and OR along with CCAM module(details in suppl. material), mostly giving incremental improvements, proving utility of each individual part. Overall, as seen in Tab. 1, we achieve improvements of about 2.63%, 2.27%, 1.77% for semantic metrics for 372, 744, and 2,975 samples of Cityscapes dataset respectively, and depth metrics by 5.3% in parallel. Fig. 4 highlights qualitative improvement seen for both depth and semantic segmentation network. We took a closer look to narrow down the classes, which are most positively impacted by our training strategy. We find that much of improvement is mainly seen for low-mIoU-classes such as motorcycle, wall, rider, and movable-classes such as bicycle, train, truck, and bus. saturated-class such as building, vegetation, car, sky, and road shows marginal improvements as these classes already have achieved about 90% mIoU (Qualitative details in supplementary material).

ScanNet Dataset At last, as part of verifying generalizibility of our proposed CCAM module, we conduct experiments on ScanNet [10], which is an indoor dataset, entirely different from Cityscapes dataset in more than one aspects. We observe that with no obvious changes to our training strategy, simply employing CCAM module, our model improves semantics metrics by 2.52% and depth metrics by 5.7% as shown in Tab. 4. With more than 3% mIoU improvement for classes such as *door, box, screen,* and *cabinet*. Due to dearth of space, we provide more details about data pre-processing and structuring, train/val split, and class wise mIoU numbers in the suppl. material.

5. Conclusions

Through this work, we go on to establish, how effective transfer of features between semantics and depth modules, could result in substantial performance gain for both the tasks, in a semi-supervised setting. We follow this up with an intelligent and diverse data augmentation for both depth and semantics. We hope these encouraging results would further push the research community in working towards finding much more efficient and effective ways for multi-task learning.

References

- [1] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep cnns? *arXiv preprint arXiv:1810.09102*, 2018.
- [2] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017.
- [4] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587, 2017.
- [5] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 3640– 3649, 2016.
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 801–818, 2018.
- [7] Po-Yi Chen, Alexander H. Liu, Yen-Cheng Liu, and Yu-Chiang Frank Wang. Towards scene understanding: Un-supervised monocular depth estimation with semantic-aware representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.
- [8] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, pages 794–803. PMLR, 2018.
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [10] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In Proc. Computer Vision and Pattern Recognition (CVPR), IEEE, 2017.
- [11] Jifeng Dai, Kaiming He, and Jian Sun. Convolutional feature masking for joint object and stuff segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3992–4000, 2015.
- [12] Tom van Dijk and Guido de Croon. How do neural networks see depth in single images? In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 2183– 2191, 2019.
- [13] Clement Farabet, Camille Couprie, Laurent Najman, and Yann LeCun. Learning hierarchical features for scene la-

beling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1915–1929, 2012.

- [14] Zhengyang Feng, Qianyu Zhou, Qiqi Gu, Xin Tan, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Dmt: Dynamic mutual training for semi-supervised learning. arXiv preprint arXiv:2004.08514, 2020.
- [15] Geoff French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations. arXiv preprint arXiv:1906.01916, 2019.
- [16] Yuan Gao, Jiayi Ma, Mingbo Zhao, Wei Liu, and Alan L Yuille. Nddr-cnn: Layerwise feature fusing in multi-task cnns by neural discriminative dimensionality reduction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3205–3214, 2019.
- [17] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision*, pages 740–756. Springer, 2016.
- [18] James J Gibson. The perception of the visual world. 1950.
- [19] Clément Godard, Oisin Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with leftright consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 270–279, 2017.
- [20] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3828–3838, 2019.
- [21] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. arXiv preprint arXiv:2002.12319, 2020.
- [22] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multitask learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 270–287, 2018.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016.
- [24] Falk Heuer, Sven Mantowsky, Saqib Bukhari, and Georg Schneider. Multitask-centernet (mcn): Efficient and diverse multitask learning using an anchor free approach. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 997–1005, 2021.
- [25] Lukas Hoyer, Dengxin Dai, Yuhua Chen, Adrian Koring, Suman Saha, and Luc Van Gool. Three ways to improve semantic segmentation with self-supervised depth estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11130–11140, 2021.
- [26] Wei-Chih Hung, Yi-Hsuan Tsai, Yan-Ting Liou, Yen-Yu Lin, and Ming-Hsuan Yang. Adversarial learning for semi-supervised semantic segmentation. arXiv preprint arXiv:1802.07934, 2018.

- [27] Pan Ji, Runze Li, Bir Bhanu, and Yi Xu. Monoindoor: Towards good practice of self-supervised monocular depth estimation for indoor environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12787–12796, 2021.
- [28] Pan Ji, Qingan Yan, Yuxin Ma, and Yi Xu. Georefine: Selfsupervised online depth refinement for accurate dense mapping. arXiv preprint arXiv:2205.01656, 2022.
- [29] Jianbo Jiao, Ying Cao, Yibing Song, and Rynson Lau. Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In *Proceedings of the European Conference on Computer Vision*, September 2018.
- [30] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 7482– 7491, 2018.
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [32] Marvin Klingner, Jan-Aike Termöhlen, Jonas Mikolajczyk, and Tim Fingscheidt. Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In *European Conference on Computer Vision*, pages 582–600. Springer, 2020.
- [33] Iasonas Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6129–6138, 2017.
- [34] Samuli Laine and Timo Aila. Temporal ensembling for semisupervised learning. arXiv preprint arXiv:1610.02242, 2016.
- [35] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [36] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5267–5276, 2019.
- [37] Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. arXiv preprint arXiv:1805.10180, 2018.
- [38] Kunpeng Li, Ziyan Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 9215–9223, 2018.
- [39] Shuai Li, Kui Jia, Yuxin Wen, Tongliang Liu, and Dacheng Tao. Orthogonal deep neural networks. *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 43(4):1352– 1368, 2021.
- [40] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for highresolution semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1925–1934, 2017.

- [41] Guosheng Lin, Chunhua Shen, Anton Van Den Hengel, and Ian Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3194–3203, 2016.
- [42] Jiachen Liu, Pan Ji, Nitin Bansal, Changjiang Cai, Qingan Yan, Xiaolei Huang, and Yi Xu. Planemvs: 3d plane reconstruction from multi-view stereo. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8665–8675, 2022.
- [43] Shikun Liu, Edward Johns, and Andrew J Davison. Endto-end multi-task learning with attention. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1871–1880, 2019.
- [44] Shikun Liu, Edward Johns, and Andrew J Davison. Endto-end multi-task learning with attention. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 1871–1880, 2019.
- [45] Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. arXiv preprint arXiv:1506.04579, 2015.
- [46] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [47] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Philip S Yu. Learning multiple tasks with multilinear relationship networks. *Advances in neural information processing systems*, 30, 2017.
- [48] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1851–1860, 2019.
- [49] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 3994–4003, 2016.
- [50] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3994–4003, 2016.
- [51] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and lowlevel consistency. *IEEE Transactions on Pattern Analysis* and Machine Intelligence, 2019.
- [52] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993, 2018.
- [53] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In Advances in Neural Information Processing Systems, pages 2204–2212, 2014.
- [54] Viktor Olsson, Wilhelm Tranheden, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 1369–1378, 2021.

- [55] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semisupervised semantic segmentation with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12674– 12684, 2020.
- [56] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017.
- [57] Pierluigi Zama Ramirez, Matteo Poggi, Fabio Tosi, Stefano Mattoccia, and Luigi Di Stefano. Geometry meets semantics for semi-supervised monocular depth estimation. In Asian Conference on Computer Vision, pages 298–313. Springer, 2018.
- [58] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. Advances in neural information processing systems, 30, 2017.
- [59] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [60] Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. Latent multi-task architecture learning. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pages 4822–4829, 2019.
- [61] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [62] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. Advances in neural information processing systems, 31, 2018.
- [63] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. arXiv preprint arXiv:1312.6229, 2013.
- [64] Yusuke Shinohara. Adversarial Multi-Task Learning of Deep Neural Networks for Robust Speech Recognition. In Proc. Interspeech 2016, pages 2369–2372, 2016.
- [65] Ayan Sinha, Zhao Chen, Vijay Badrinarayanan, and Andrew Rabinovich. Gradient adversarial training of neural networks. arXiv preprint arXiv:1806.08028, 2018.
- [66] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *International Conference on Machine Learning*, pages 9120–9132. PMLR, 2020.
- [67] Ximeng Sun, Rameswar Panda, Rogerio Feris, and Kate Saenko. Adashare: Learning what to share for efficient deep multi-task learning. arXiv preprint arXiv:1911.12423, 2019.
- [68] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. arXiv preprint arXiv:1703.01780, 2017.

- [69] Lokender Tiwari, Pan Ji, Quoc-Huy Tran, Bingbing Zhuang, Saket Anand, and Manmohan Chandraker. Pseudo rgb-d for self-improving monocular slam and depth prediction. In *European conference on computer vision*, pages 437–455, 2020.
- [70] Jiayun Wang, Yubei Chen, Rudrasis Chakraborty, and Stella X. Yu. Orthogonal convolutional neural networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, June 2020.
- [71] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018.
- [72] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12275–12284, 2020.
- [73] Jamie Watson, Oisin Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1164–1174, 2021.
- [74] Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 7268– 7277, 2018.
- [75] John T Wixted and Sharon L Thompson-Schill. Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, Language and Thought, volume 3. John Wiley & Sons, 2018.
- [76] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. arXiv preprint arXiv:1904.12848, 2019.
- [77] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 675–684, 2018.
- [78] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- [79] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019.
- [80] Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [81] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Yan Yan, Nicu Sebe, and Jian Yang. Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4106–4115, 2019.

- [82] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2881–2890, 2017.
- [83] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 2921– 2929, 2016.
- [84] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017.
- [85] Yuliang Zou, Pan Ji, Quoc-Huy Tran, Jia-Bin Huang, and Manmohan Chandraker. Learning monocular visual odometry via self-supervised long-term modeling. In *European Conference on Computer Vision*, pages 710–727, 2020.