

Dynamic Feature Fusion for Semantic Edge Detection

Yuan Hu^{1,2}, Yunpeng Chen³, Xiang Li^{1,2} and Jiashi Feng³

¹Institute of Remote Sensing and Digital Earth, CAS, Beijing 100094, China

²University of Chinese Academy of Sciences, Beijing 100049, China

³National University of Singapore

{huyuan, lixiang01}@radi.ac.cn, chenyunpeng@u.nus.edu, elefjia@nus.edu.sg

Abstract

Features from multiple scales can greatly benefit the semantic edge detection task if they are well fused. However, the prevalent semantic edge detection methods apply a fixed weight fusion strategy where images with different semantics are forced to share the same weights, resulting in universal fusion weights for all images and locations regardless of their different semantics or local context. In this work, we propose a novel dynamic feature fusion strategy that assigns different fusion weights for different input images and locations adaptively. This is achieved by a proposed weight learner to infer proper fusion weights over multi-level features for each location of the feature map, conditioned on the specific input. In this way, the heterogeneity in contributions made by different locations of feature maps and input images can be better considered and thus help produce more accurate and sharper edge predictions. We show that our model with the novel dynamic feature fusion is superior to fixed weight fusion and also the naïve location-invariant weight fusion methods, via comprehensive experiments on benchmarks Cityscapes and SBD. In particular, our method outperforms all existing well established methods and achieves new state-of-the-art.

1 Introduction

The task of semantic edge detection (SED) is aimed at both detecting visually salient edges and recognizing their categories, or more concretely, locating fine edges utilizing low-level features and meanwhile identifying semantic categories with abstracted high-level features. An intuitive way for a deep CNN model to achieve both targets is to integrate high-level semantic features with low-level category-agnostic edge features via a fusion model, which is conventionally designed following a fixed weight fusion strategy, independent of the input, as illustrated in the top row in Figure 1.

In many existing deep SED models [Yu *et al.*, 2017; Liu *et al.*, 2018b; Yu *et al.*, 2018], fixed weight fusion of multi-level features is implemented through 1×1 convolution, where the learned convolution kernel serves as the fusion weights. However, this fusion strategy cannot fully

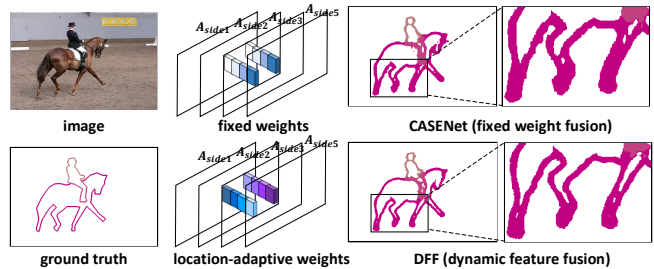


Figure 1: Illustration of fixed weight fusion and proposed dynamic feature fusion (DFF). The proposed model actively learns customized fusion weights for each location, to fuse low-level features ($A_{side1} \sim A_{side3}$) and a high-level feature (A_{side5}). In this way, sharper boundaries are produced through dynamic feature fusion compared with CASENet (fixed weight fusion model).

exploit multi-level information, especially the low-level features. This is because, first, it applies the same fusion weights to all the input images and ignores their variations in contents, illumination, *etc.* The distinct properties of a specific input need be treated adaptively for revealing the subtle edge details. Besides, for a same input image, different spatial locations on the corresponding feature map convey different information, but the fixed weight fusion manner applies the same weights to all these locations, regardless of their different semantic classes or object parts. This would unfavorably force the model to learn universal fusion weights for all the categories and locations. Consequently, a bias would be caused toward high-level features, and the power of multi-level response fusion is significantly weakened.

In this work, we propose a *Dynamic Feature Fusion* (DFF) method that assigns adaptive fusion weights to each location individually, aiming to generate a fused edge map adaptive to the specific content of each image, as illustrated in the bottom row in Figure 1. In particular, we design a novel location-adaptive weight learner that actively learns customized location-specific fusion weights conditional on the feature map content for multi-level response maps. As shown in Figure 1, low-level features ($A_{side1} \sim A_{side3}$) and a high-level feature (A_{side5}) are merged to produce the final fused output. Low-level feature maps give high response on fine details, such as the edges inside objects, whereas high-level ones are coarser and only exhibit strong response at the object-level boundaries. This location-adaptive weight

learner tailors fusion weights for each individual location. For example, for the boundaries of the horse, fusion weights are biased toward low-level features to fully take advantage of the accurate located edges. For the interior of the horse, higher weights are assigned to high-level features to suppress the fragmentary and trivial edge responses inside the object.

The proposed DFF model consists of two main components: a *feature extractor with a normalizer* and an *adaptive weight fusion module*. The feature extractor primarily scales the multi-level responses to the same magnitude, preparing for the down-streaming fusion operation. The adaptive weight fusion module performs following two computations. First, it dynamically generates location-specific fusion weights conditioned on the image content. Then, the location-specific fusion weights are applied to actively fuse the high-level and low-level response maps. The adaptive weight fusion module is capable of fully excavating the potentialities of multi-level responses, especially the low-level ones, to produce better fusion output for every single location.

In summary, our main contributions are:

- For the first time, this work reveals limitations of the popular fixed weight fusion for SED, and explains why it does not produce satisfactory fusion results as expected.
- We propose a dynamic feature fusion (DFF) model. To our best knowledge, it is the first work to learn adaptive fusion weights conditioned on input contents to merge multi-level features in the research field of SED.
- The proposed DFF model achieves new state-of-the-art on the SED task.

2 Related Work

The task of category-aware semantic edge detection was first introduced by [Hariharan *et al.*, 2011]. It is tackled usually as a multi-class problem [Hariharan *et al.*, 2011; Bertasius *et al.*, 2015; 2016; Maninis *et al.*, 2016] at first, in which only one semantic class is associated with each located boundary pixel. From CASENet [Yu *et al.*, 2017], researchers begin to address this task as a multi-label problem where each edge pixel can be associated with more than one semantic class simultaneously. Recently, deep learning based models, such as SEAL [Yu *et al.*, 2018] and DDS [Liu *et al.*, 2018b], further lift the performance of semantic edge detection to new state-of-the-art.

The tradition of employing fixed weight fusion can be pinpointed to HED [Xie and Tu, 2015], in which a weighted-fusion layer (implemented by 1×1 convolution layer) is employed to merge the side-outputs. RCF [Liu *et al.*, 2017] and RDS [Liu and Lew, 2016] follow this simple strategy to perform category-agnostic edge detection. CASENet [Yu *et al.*, 2017], SEAL [Yu *et al.*, 2018] and DDS [Liu *et al.*, 2018b] extend the approach with a K -grouped 1×1 convolution layer to generate the K -channel fused activation map. In this paper, we demonstrate the above fixed fusion strategy cannot sufficiently leverage multi-scale responses for producing better fused output. Instead, our proposed adaptive weight fusion module enables the network to actively learn the location-

aware fusion weights conditioned on the individual input content.

Multi-level representations within a convolutional neural network have been shown effective in many vision tasks, such as object detection [Liu *et al.*, 2016; Lin *et al.*, 2017], semantic segmentation [Long *et al.*, 2015; Yu and Koltun, 2015], boundary detection [Xie and Tu, 2015; Liu *et al.*, 2017], object symmetry detection [Ke *et al.*, 2017; Liu *et al.*, 2018a]. Responses from different stages of a network tend to vary significantly in magnitude. Scaling the multi-scale activations to a similar magnitude would benefit the following prediction or fusion consistently. For example, SSD [Liu *et al.*, 2016] performs L2 normalization on low-level feature maps before making multi-scale predictions. [Li *et al.*, 2018] adds a scaling layer for learning a fusion scale to combine the channel attention and spatial attention outputs. In this paper, we adopt a feature extractor with a normalizer for all multi-level activations to deal with the bias.

Our proposed method is also related with the dynamic filter networks [Jia *et al.*, 2016] in which sample-specific filter parameters are generated based on the input. The filters with a specific kernel size are dynamically produced to enable local spatial transformations on the input feature map, which serve as dynamic convolution kernels. Comparatively, in our method, the adaptive fusion weights are applied on the multi-level response maps to obtain a desired fused output and a feature extractor with a normalizer is proposed to alleviate the bias during training.

3 Feature Fusion with Fixed Weights

Before expounding the proposed model, we first introduce the notations used in our model and revisit the fixed weight fusion in CASENet [Yu *et al.*, 2017] as preliminaries.

CASENet adopts ResNet-101 with some minor modifications to extract multi-level features. Based on the backbone, a 1×1 convolution layer and a following upsampling layer are connected to the output of the first three and the top stacks of residual blocks, producing three single channel feature maps $\{A_{side1}, A_{side2}, A_{side3}\}$ and a K -channel class activation map $A_{side5} = \{A_{side5}^1, \dots, A_{side5}^K\}$ respectively. Here K is the number of categories. The shared concatenation replicates the bottom features $\{A_{side1}, A_{side2}, A_{side3}\}$ for K times to separately concatenate with each of the K top activations in A_{side5} :

$$A_{cat} = \{A_{cat}^1, \dots, A_{cat}^K\}, \quad (1)$$

$$A_{cat}^i = \{A_{side5}^i, A_{side1}, A_{side2}, A_{side3}\}, i \in [1, K]. \quad (2)$$

The resulting concatenated activation map A_{cat} is then fed into the K -grouped 1×1 conv layer to produce the fused activation map with K channels:

$$A_{fuse}^i = w_1^i A_{side5}^i + w_2^i A_{side1} + w_3^i A_{side2} + w_4^i A_{side3} \quad (3)$$

$$A_{fuse} = \{A_{fuse}^i\}, i \in [1, K] \quad (4)$$

where $(w_1^i, w_2^i, w_3^i, w_4^i)$ are the parameters of the K -grouped 1×1 conv layer, serving as fusion weights for the i th category. We omit the bias term in Eqn. (3) for simplicity. One can refer to [Yu *et al.*, 2017] for more details.

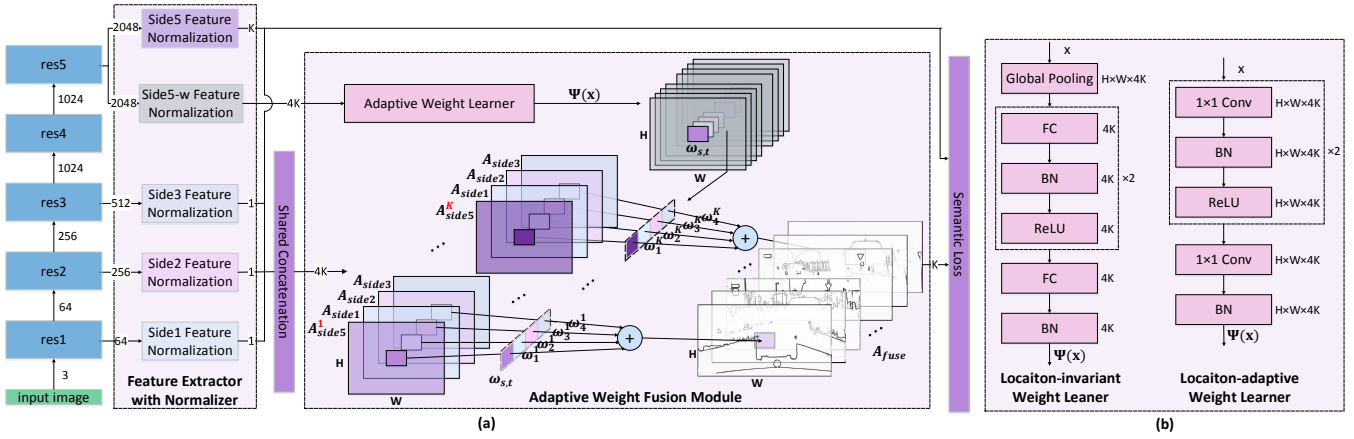


Figure 2: Overall architecture of our model. (a) The input image is fed into the ResNet backbone to generate a set of features (response maps) with different scales. Side feature normalization blocks are connected to the first three and the fifth stack of residual blocks to produce Side1-3 and Side5 response maps with the same response magnitude. Shared concatenation (Eqn. (1) and (2)) is then applied to concatenate Side1-3 and Side5. The *Side5-w feature normalization* block followed by a *location-adaptive weight learner* form another branch extended from res5 to predict dynamic location-aware fusion weights $\Psi(x)$. Then element-wise multiplication and category-wise summation are applied to the location-aware fusion weights $\Psi(x)$ and the concatenated response map A_{cat} to generate the final fused output A_{fuse} . The semantic loss is employed to supervise A_{side5} and the final fused output A_{fuse} . (b) *Location-invariant weight learner* and *location-adaptive weight learner* take as input the feature map x and output location-invariant fusion weights and location-adaptive fusion weights $\Psi(x)$ respectively.

CASNet imposes the same fixed set of fusion weights ($w_1^i, w_2^i, w_3^i, w_4^i$) of the i th category upon all the images across all the locations of the feature maps. By reproducing CASNet, we make following observations. i) The weight w_1^i always surpasses the other weights (w_2^i, w_3^i, w_4^i) for each class significantly. ii) By evaluating the performance of A_{side5} and the fused output A_{fuse} , we find they give almost the same performance. This implies that the low-level feature responses $\{A_{side1}, A_{side2}, A_{side3}\}$ contribute very little to the final fused output although they contain fine-grained edge location and structure information. The final decision is mainly determined by high-level features. As such, the CASNet cannot exploit low-level information sufficiently and produce only coarse edges.

4 Proposed Model

To address the above limitations of existing *Semantic Edge Detection* (SED) models in fusing features with fixed weights, we develop a new SED model with dynamic feature fusion. The proposed model fuses multi-level features through two modules: 1) the feature extractor with a normalizer is to normalize the magnitude scales of multi-level features and 2) the adaptive weight fusion module is to learn adaptive fusion weights for different locations of multi-level feature maps (see Figure 2 (a)).

4.1 Dynamic Feature Fusion

The first module is inspired by following observations. In CASNet [Yu *et al.*, 2017], the edge responses from the top layers are much stronger than the ones from other three bottom outputs. Such variation in scales of activations biases the multi-level feature fusion to the responses from the top layers. Moreover, the top-layer output is much similar to the ground-truth for the training examples, due to the direct supervision. Applying the same weights to all the locations forces the net-

work to learn much higher fusion weights for the top-level features, undesirably ignoring contributions from low-level features. This further inhibits the low-level features from providing fine edge information for detecting object boundaries in the cases of heavily biased distributions of edge and non-edge pixels.

Therefore, before feature fusion, we first deal with the scale variation of multi-level responses by normalizing their magnitudes. In this way, the following adaptive weight learner can get rid of distractions from scale variation and learn effective fusion weights more easily.

A feature extractor with a normalizer (see Figure 2 (a)) is employed to normalize multi-level responses to the similar magnitude. Concretely, side feature normalization blocks in the module are responsible for normalizing feature maps of the corresponding level.

To achieve our proposed dynamic feature fusion, we devise two different schemes for predicting the adaptive fusion weights, *i.e.*, location-invariant and location-adaptive fusion weights (see Figure 2 (b)). The former treats all locations in the feature maps equally and universal fusion weights are learned according to specific input adaptively. The later one adjusts the fusion weights conditioned on location features of the image and lifts the contributions of low-level features to locating fine edges along object boundaries.

Concretely, given the activation of the multi-level side outputs A_{side} of size $H \times W$, we aim to obtain a fused output $A_{fuse} = f(A_{side})$ by aggregating multi-level responses. The fusion operation in CASNet in Eqn. (3) can be written as the following generalized form:

$$A_{fuse} = f(A_{side}; \mathbf{W}) \quad (5)$$

where f encapsulates operations in Eqn. (3) and $\mathbf{W} = (w_1^i, w_2^i, w_3^i, w_4^i)$ denotes the fusion weights.

Differently, we propose an adaptive weight learner to actively learn the fusion weights conditioned on the feature map

itself, which is formulated as

$$A_{fuse} = f(A_{side}; \Psi(x)), \tag{6}$$

where x denotes the feature map. The above formulations characterize the essential difference between our proposed adaptive weight fusion method and the fixed weight fusion method. We enforce the fusion weights \mathbf{W} to explicitly depend on the feature map x , i.e. $\mathbf{W} = \Psi(x)$. Different input image would produce different feature maps x , thus induce different parameters $\Psi(x)$, and then lead to modifications to the adaptive weight learner $f(\cdot; \cdot)$ dynamically. In this way, the semantic edge detection model can fast adapt to the input image and favourably learn proper multi-level response fusion weights in an end-to-end manner.

Regarding $\Psi(x)$, corresponding to the two fusion weight schemes, there are two adaptive weight learners, namely location-invariant weight learner and location-adaptive weight learner. The location-invariant weight learner learns $4K$ fusion weights in total as shown in the following equation, which are shared by all locations of feature maps to be fused:

$$\Psi(x) = (w_1^i, w_2^i, w_3^i, w_4^i), i \in [1, K]. \tag{7}$$

However, the location-adaptive weight learner generates $4K$ fusion weights for each spatial location, which results in $4KH$ weighting parameters in total.

$$\Psi(x) = (\mathbf{w}_{s,t}), s \in [1, H], t \in [1, W] \tag{8}$$

$$\mathbf{w}_{s,t} = ((w_1^i)_{s,t}, (w_2^i)_{s,t}, (w_3^i)_{s,t}, (w_4^i)_{s,t}), i \in [1, K] \tag{9}$$

The location-invariant weight learner generates universal fusion weights for all locations, while the location-adaptive weight learner tailors fusion weights for each location considering the spatial variety.

4.2 Network Architecture

Our proposed network architecture is based on ResNet [He *et al.*, 2016] and adopts the same modifications as CASNet [Yu *et al.*, 2017] to preserve low-level edge information, as shown in Figure 2. Side feature normalization blocks are connected to the first three and the fifth stack of residual blocks. This block consists of a 1×1 convolution layer, a Batch Normalization (BN) [Ioffe and Szegedy, 2015] layer and a deconvolution layer. The 1×1 convolution layer produces single and K channel response maps for Side1-3 and Side5 respectively. The BN layer is applied on the output of the 1×1 convolution layer to normalize the multi-level responses to the same magnitude. Deconvolution layers are then used to upsample the response maps to the original image size.

Another side feature normalization block is connected to the fifth stack of the residual block, in which a $4K$ -channel feature map is produced. The adaptive weight learner then takes in the output of the Side5-w feature normalization block to predict the dynamic fusion weights $\Psi(x)$. We design two instantiations for location-invariant weight learner and location-adaptive weight learner respectively, as detailed below.

Method	Normalizer	Learner			mean	Δ mean	run time (s)
		invariant	adaptive	softmax			
Baseline					78.4		0.86
Ours	✓				78.7	+ 0.3	0.86
	✓	✓			79.1	+ 0.7	1.27
	✓		✓		80.1	+ 1.7	1.00
	✓			✓	80.4	+ 2.0	0.99
					78.9	+ 0.5	0.99

Table 1: The effectiveness of each component in our proposed model on Cityscapes dataset. The mean value of MF scores over all categories is measured by %. Run time is tested on 1 GPU with batch size 1.

Location-invariant weight learner. Location-invariant weight learner (see Figure 2 (b)) is a naive version of adaptive weight learner. The output of the Side5-w feature normalization block x with size $H \times W \times 4K$ is taken as the input of a global average pooling layer to produce a $4K$ -channel vector. After that, a block of alternatively connected FC, BN and ReLU layers is repeated twice and followed by FC and BN layers to generate the $4K$ location-invariant fusion weights for the fused response map A_{fuse} with size $H \times W \times 4K$. These fusion weights are conditional on the input content but all locations share the same fusion parameters.

Location-adaptive weight learner. We further propose the location-adaptive weight learner (see Figure 2 (b)) to solve the shortcomings of fixed weight fusion. The feature map x is forwarded to a block of alternatively connected 1×1 conv, BN and ReLU layers, repeated twice and followed by a 1×1 conv and BN layer, to generate location-aware fusion weights $\Psi(x)$ with size $H \times W \times 4K$. Element-wise multiplication and category-wise summation are then applied to the location-aware fusion weights and the fused response map A_{fuse} to generate the final fused output. The location-adaptive weight learner predicts $4K$ fusion weights for each spatial location on Side1-3 and Side5 to allow low-level responses to provide fine edge locations for object boundaries. One may consider applying softmax activation function to the output of the adaptive weight learner to learn fusion weights mutually-exclusive for the side feature maps to be fused respectively. The ablation experiments in Section 5.2 demonstrate that the activation function hampers the performance in the case of generating adaptive fusion weights.

5 Experiments

5.1 Experimental Setup

Datasets and augmentations. We evaluate our proposed method on two popular benchmarks for semantic edge detection: Cityscapes [Cordts *et al.*, 2016] and SBD [Hariharan *et al.*, 2011]. During training, we use random mirroring and random cropping on both Cityscapes and SBD. For Cityscapes, we also augment each training sample with random scaling in $[0.75, 2]$, and the base size for random scaling is set as 640. For SBD, we augment the training data by resizing each image with scaling factors $\{0.5, 0.75, 1, 1.25, 1.5\}$ following [Yu *et al.*, 2017]. During testing, the images with original image size are used for Cityscapes and the images are padded to 512×512 for SBD.

Training strategies. The proposed network is built on ResNet [He *et al.*, 2016] pretrained on ImageNet [Deng *et al.*

Method	backbone	crop size	road	side.	buil.	wall	fence	pole	light	sign	vege.	terrain	sky	person	rider	car	truck	bus	train	mot.	bike	mean
DFF	ResNet50	512	93.1	86.5	89.0	58.8	63.5	84.8	83.7	83	91.8	71.9	92.2	90.7	81.3	94.7	58.5	72.7	50.8	71.7	86.6	79.2
	ResNet50	640	93.3	86.5	89.5	58.2	64.3	85.2	85.5	84.3	92.2	72.0	92.6	91.2	82.6	94.9	60.0	77.2	56.9	74.5	86.3	80.4
	ResNet101	512	93.2	86.3	89.7	60.0	64.3	86.0	85.2	83.7	92.0	71.3	92.3	91.1	83.4	94.8	64.5	77.4	59.9	72.1	86.9	80.7

Table 2: Results of the proposed DFF using different backbone networks and crop sizes. All MF scores are measured by %

Method	road	side.	buil.	wall	fence	pole	light	sign	vege.	terrain	sky	person	rider	car	truck	bus	train	mot.	bike	mean
CASENet	86.6	78.8	85.1	51.5	58.9	70.1	70.8	74.6	83.5	62.9	79.4	81.5	71.3	86.9	50.4	69.5	52.0	61.3	80.2	71.3
DDS-R	90.5	84.2	86.2	57.7	61.4	85.1	83.8	80.4	88.5	67.6	88.2	89.9	80.1	91.8	58.6	76.3	56.2	68.8	87.3	78.0
DFF	93.2	86.3	89.7	60.0	64.3	86.0	85.2	83.7	92.0	71.3	92.3	91.1	83.4	94.8	64.5	77.4	59.9	72.1	86.9	80.7
SEAL (0.0035)	87.6	77.5	75.9	47.6	46.3	75.5	71.2	75.4	80.9	60.1	87.4	81.5	68.9	88.9	50.2	67.8	44.1	52.7	73.0	69.1
DFF (0.0035)	89.4	80.1	79.6	51.3	54.5	81.3	81.3	81.2	83.6	62.9	89.0	85.4	75.8	91.6	54.9	73.9	51.9	64.3	76.4	74.1

Table 3: Comparison with state-of-the-arts on Cityscapes dataset. All MF scores are measured by %

al., 2009]. The network is trained with standard reweighted cross-entropy loss [Yu et al., 2017] and optimized by SGD using PyTorch [Paszke et al., 2017]¹. We set the base learning rate to 0.08 and 0.05 for Cityscapes and SBD respectively; the “poly” policy is used for learning rate decay. The crop size, batch size, training epoch, momentum, weight decay are set to $640 \times 640 / 352 \times 352, 8 / 16, 200 / 10, 0.9, 1e - 4$ for Cityscapes and SBD respectively. All experiments are performed using 4 NVIDIA TITAN Xp(12GB) GPUs with synchronized batch normalization [Zhang et al., 2018]. We use the same data augmentation and hyper-parameters for reproducing CASENet and training the proposed model for fair comparison.

Evaluation protocol. We follow [Yu et al., 2018] to evaluate the performance with stricter rules than the benchmark used in [Yu et al., 2017; Liu et al., 2018b]. The ground truth maps are downsampled into half size of original dimensions for Cityscapes, and are generated with instance-sensitive edges for both datasets. Maximum F-measure (MF) at optimal dataset scale (ODS) with matching distance tolerance set as 0.02 is used for Cityscapes and SBD evaluation following our baseline methods. For fair comparison with [Yu et al., 2018], we also set the distance tolerance to 0.0035 for Cityscapes dataset.

5.2 Ablation Study

We first conduct ablation experiments on Cityscapes to investigate the effectiveness of each newly proposed module. We use CASENet as the baseline model and add our proposed normalizer and learner on it. All the ablation experiments use the same settings with 640×640 crop size and ResNet50 backbone. We report the reproduced baseline CASENet trained with the same training settings for fair comparison, which actually has higher accuracy than the original paper. Results are summarized in Table 1.

Normalizer. We first evaluate the effect of the proposed normalizer in the feature extractor. As shown in Table 1 (row 2), the normalizer provides 0.3% performance gain compared with the baseline which does not normalize the multi-level features to the same magnitude before fusion. It shows that normalized features are more suitable for multi-level information fusion than the unnormalized ones. This is probably because the unnormalized features have dramatically differ-

ent magnitudes, making the learning process biased to the features with higher magnitude.

Location-invariant learner. Row 3 in Table 1 shows the results of further adding the proposed location-invariant learner which adaptively fuses multi-level features conditioned on the input image. As shown in Table 1, the location-invariant learner provides another 0.4% performance gain compared with *baseline+normalizer* that assigns fixed fusion weights without considering the uniqueness of each input image. The performance boost successfully shows that different input images with different illumination, semantics, etc., may require different fusion weights and applying universal weights to all training images will lead to less satisfactory results.

Location-adaptive learner. The location-adaptive learner provides location-aware fusion weights that are conditioned on not only the input image but also the local context, which further improves the location-invariant learner. As shown in Table 1 row 3 and row 5, location-adaptive learner improves the performance significantly from 79.1% to 80.4%, indicating a fusion module should consider both the uniqueness of each input image and the local context at each location. The proposed fusion module in DFF successfully captures both of them and thus delivers remarkable performance gain. Besides, it can be seen that the location-adaptive learner heavily relies on pre-normalized multi-level features, since the performance degrades by 1.5% when the proposed normalizer is removed from the framework, as shown in row 5 and row 6. This further justifies the effectiveness of the proposed normalizer for rescaling the activation values to the same magnitude before conducting fusion.

Activation function. Row 4 and row 5 in Table 1 show the cases of further constraining the generated fusion weights for better performance. In particular, we append a softmax layer to the end of location-adaptive weight learner (Figure 2 (b)) to force the learned fusion weights for each category to add up to 1, and compare its results with the unconstrained version. However, we observe performance drop of 0.3% with softmax activation function adopted. Our intuitive explanation is that applying softmax activation function to the fusion weights normalizes their values to the range $[0, 1]$, which may actually be negative for some response maps or some locations. Furthermore, adopting softmax activation function enforces the fusion weights of multi-level response maps for each category to be mutually exclusive, which makes the

¹Codes will be released soon.

Method	aer.	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mot.	per.	pot.	sheep	sofa	train	tv	mean
CASENet	83.6	75.3	82.3	63.1	70.5	83.5	76.5	82.6	56.8	76.3	47.5	80.8	80.9	75.6	80.7	54.1	77.7	52.3	77.9	68.0	72.3
DDS-R	85.4	78.3	83.3	65.6	71.4	83.0	75.5	81.3	59.1	75.7	50.7	80.2	82.7	77.0	81.6	58.2	79.5	50.2	76.5	71.2	73.3
SEAL	84.5	76.5	83.7	64.9	71.7	83.8	78.1	85.0	58.8	76.6	50.9	82.4	82.2	77.1	83.0	55.1	78.4	54.4	79.3	69.6	73.8
DFF	86.5	79.5	85.5	69.0	73.9	86.1	80.3	85.3	58.5	80.1	47.3	82.5	85.7	78.5	83.4	57.9	81.2	53.0	81.4	71.6	75.4

Table 4: Comparison with state-of-the-arts on SBD dataset. All MF scores are measured by %

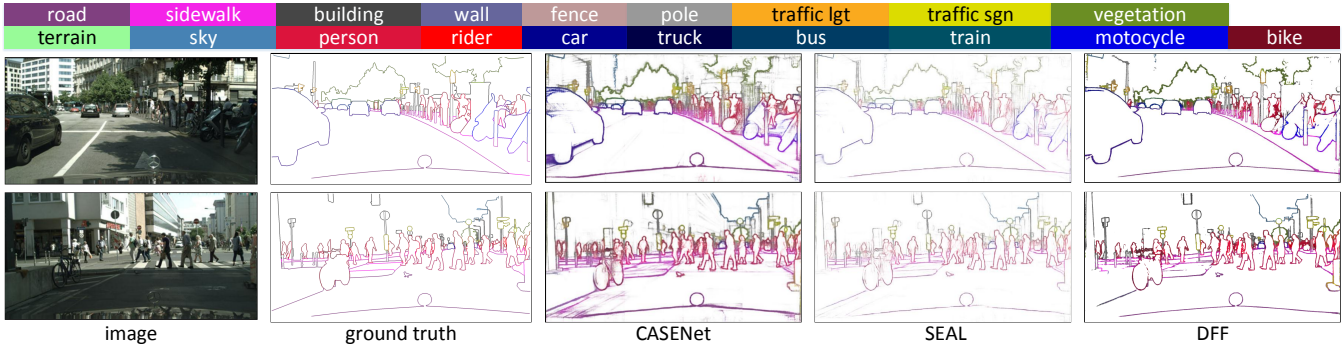


Figure 3: Qualitative comparison on Cityscapes among ground truth, CASENet, SEAL and DFF (ordering from left to right in the figure). Best viewed in color.

response maps compete with each other to be more important. However, it is desirable that each response map be given enough freedom to be active and be useful to the final fused output. Therefore, we do not use any activation functions to constrain the fusion weights in the proposed DFF model.

Deeper backbone network. We also perform experiments on deeper backbones to verify our model’s generalization. As shown in Table 2, we evaluate our model performance on ResNet101 with 512×512 crop size. Note that we do not use 640×640 , simply because of the GPU memory limitation. In order to make a fair comparison, we also train a DFF on ResNet50 with 512×512 crop size. The result shows DFF also generalizes well on deep ResNet101 and outperforms its counterpart (ResNet50, 512×512) by 1.5%. Higher performance might be achieved if we can train the ResNet101 with 640×640 crop size using GPUs with larger memory, *e.g.* Nvidia V100 (32GB).

5.3 Comparison with State-of-the-arts

We compare the performance of the proposed DFF model with state-of-the-art semantic edge detection methods [Yu *et al.*, 2017; 2018; Liu *et al.*, 2018b] on Cityscapes and SBD. All of them are build on ResNet-101 [He *et al.*, 2016].

Results on Cityscapes

We compare the proposed DFF model with CASENet and DDS, and evaluate MF (ODS) with matching distance tolerance set as 0.02. In addition, the matching distance tolerance is decreased to 0.0035 to make fair comparison with SEAL. As can be seen in Table 3, our DFF model outperforms all these well established baselines and achieves new state-of-the-art 80.7% MF (ODS). The DFF model is superior to other models for most classes. Specifically, the MF (ODS) of the proposed DFF model is 9.4% higher than CASENet and 2.7% higher than DDS on average. It is also worth noting that DFF achieves 5% higher than SEAL under stricter matching distance tolerance (0.0035), which reflects the DFF model can

locate more accurate edges by taking full advantage of multi-level response maps, especially the low-level ones, with dynamic feature fusion. We also visualize some prediction results for qualitative comparison shown in Figure 3. Through comparing with CASENet and SEAL, we can observe DFF can predict more accurate and clearer edges on object boundaries while suppressing the fragmentary and trivial edge responses inside the objects.

Results on SBD

Compared with Cityscapes, SBD dataset has fewer objects and especially fewer overlapping ones in each image. In addition, it contains many misaligned annotations. We compare the proposed DFF model with state-of-the-arts on SBD validation set in Table 4. The proposed DFF model outperforms all the baselines and achieves new state-of-the-art 75.4% MF (ODS), which well confirms its superiority.

6 Conclusion

In this paper, we propose a novel dynamic feature fusion (DFF) model to exploit multi-level responses for producing a desired fused output. The newly introduced normalizer scales multi-level side activations to the same magnitude preparing for the down-streaming fusion operation. Location-adaptive weight learner is proposed to actively learn customized fusion weights conditioned on the individual sample content for each spatial location. Comprehensive experiments on Cityscapes and SBD demonstrate the proposed DFF model can improve the performance by locating more accurate edges on object boundaries as well as suppressing trivial edge responses inside the objects. In the future, we will continue to improve the location-adaptive weight learner by considering both the high-level and low-level feature maps from the backbone.

Acknowledgments

Jiashi Feng was partially supported by NUS IDS R-263-000-C67-646, ECRA R-263-000-C87-133 and MOE Tier-II R-263-000-D17-112.

References

- [Bertasius *et al.*, 2015] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. High-for-low and low-for-high: Efficient boundary detection from deep object features and its applications to high-level vision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 504–512, 2015.
- [Bertasius *et al.*, 2016] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Semantic segmentation with boundary neural fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3602–3610, 2016.
- [Cordts *et al.*, 2016] 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.
- [Deng *et al.*, 2009] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [Hariharan *et al.*, 2011] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. 2011.
- [He *et al.*, 2016] 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.
- [Ioffe and Szegedy, 2015] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [Jia *et al.*, 2016] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc V Gool. Dynamic filter networks. In *Advances in Neural Information Processing Systems*, pages 667–675, 2016.
- [Ke *et al.*, 2017] Wei Ke, Jie Chen, Jianbin Jiao, Guoying Zhao, and Qixiang Ye. Srn: side-output residual network for object symmetry detection in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1068–1076, 2017.
- [Li *et al.*, 2018] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In *CVPR*, volume 1, page 2, 2018.
- [Lin *et al.*, 2017] Tsung-Yi Lin, Piotr Dollár, Ross B Girshick, Kaiming He, Bharath Hariharan, and Serge J Belongie. Feature pyramid networks for object detection. In *CVPR*, volume 1, page 4, 2017.
- [Liu and Lew, 2016] Yu Liu and Michael S Lew. Learning relaxed deep supervision for better edge detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 231–240, 2016.
- [Liu *et al.*, 2016] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [Liu *et al.*, 2017] Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Kai Wang, and Xiang Bai. Richer convolutional features for edge detection. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5872–5881. IEEE, 2017.
- [Liu *et al.*, 2018a] Chang Liu, Wei Ke, Fei Qin, and Qixiang Ye. Linear span network for object skeleton detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 133–148, 2018.
- [Liu *et al.*, 2018b] Yun Liu, Ming-Ming Cheng, Jia Wang, Le Zhang, Peng-Tao Jiang, and Yang Cao. Semantic edge detection with diverse deep supervision. *arXiv preprint arXiv:1804.02864*, 2018.
- [Long *et al.*, 2015] 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.
- [Maninis *et al.*, 2016] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc Van Gool. Convolutional oriented boundaries. In *European Conference on Computer Vision*, pages 580–596. Springer, 2016.
- [Paszke *et al.*, 2017] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [Xie and Tu, 2015] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*, pages 1395–1403, 2015.
- [Yu and Koltun, 2015] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- [Yu *et al.*, 2017] Zhiding Yu, Chen Feng, Ming-Yu Liu, and Srikumar Ramalingam. Casenet: Deep category-aware semantic edge detection. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA*, pages 21–26, 2017.
- [Yu *et al.*, 2018] Zhiding Yu, Weiyang Liu, Yang Zou, Chen Feng, Srikumar Ramalingam, BVK Vijaya Kumar, and Jan Kautz. Simultaneous edge alignment and learning. *arXiv preprint arXiv:1808.01992*, 2018.
- [Zhang *et al.*, 2018] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoqiang Wang, Amrbrish Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.