

Doubly Contrastive End-to-End Semantic Segmentation for Autonomous Driving under Adverse Weather

Jongoh Jeong
jojeong@rit.kaist.ac.kr
Jong-Hwan Kim
johkim@rit.kaist.ac.kr

Korea Advanced Institute of Science
and Technology (KAIST)
Daejeon, Republic of Korea

Abstract

Road scene understanding tasks have recently become crucial for self-driving vehicles. In particular, real-time semantic segmentation is indispensable for intelligent self-driving agents to recognize roadside objects in the driving area. As prior research works have primarily sought to improve the segmentation performance with computationally heavy operations, they require far significant hardware resources for both training and deployment, and thus are not suitable for real-time applications. As such, we propose a doubly contrastive approach to improve the performance of a more practical lightweight model for self-driving, specifically under adverse weather conditions such as fog, nighttime, rain and snow. Our proposed approach exploits both image- and pixel-level contrasts in an end-to-end supervised learning scheme without requiring a memory bank for global consistency or the pretraining step used in conventional contrastive methods. We validate the effectiveness of our method using SwiftNet on the ACDC dataset, where it achieves up to 1.34%p improvement in mIoU (ResNet-18 backbone) at 66.7 FPS (2048×1024 resolution) on a single RTX 3080 Mobile GPU at inference. Furthermore, we demonstrate that replacing image-level supervision with self-supervision achieves comparable performance when pre-trained with clear weather images.

1 Introduction

Road scene understanding tasks have recently been popular areas of research, gaining traction with the growth of deep learning and interests in intelligent driving agents. In particular, in the field of autonomous driving, achieving high performance in the task of semantic segmentation is key to recognizing roadside objects in the scene ahead [24, 47]. A number of existing semantic segmentation models follow the encoder-decoder structure [9, 7, 24, 58] or a multi-scale pyramidal encoder followed by upsampling [19, 28, 29, 40] to categorize each pixel into the defined set of object classes, with 2D RGB images as input.

However, as intelligent driving systems require safe and accurate perception of the surroundings in dynamically changing environments, most heavyweight models that focus primarily on the performance are not suitable for practical real-time deployment. In addition,

although various methods target semantic segmentation of the 19 *Cityscapes* [10] benchmark classes under clear, daytime weather conditions and have been useful in most *normal* road scenarios, their direct application under adverse driving conditions, or “unusual road or traffic conditions that were not known” as defined by the US Federal Motor Carrier Safety Administration [14] (e.g., fog, nighttime, rain, snow), has yet to be explored further [65]. In this light, a number of models and datasets have been proposed to boost the segmentation performance under such adverse conditions [12, 63, 64, 64] as there are 1.3 million death toll of road traffic crashes every year, and the risk of accidents in rainy weather, for example, is 70% higher than in normal conditions [8, 27]. In this approach, we seek to adapt and extend the segmentation capability of intelligent self-driving agents to such adverse conditions by exploiting a lightweight model and minimizing additional costs for ground truth labels.

To this end, we propose a doubly contrastive end-to-end learning method for a lightweight semantic segmentation model under adverse weather scenarios. Given an RGB image as input, our proposed strategy optimizes according to the supervised contrastive objective in both image- and pixel-levels in an end-to-end manner, in addition to the segmentation loss. In contrast to the previous contrastive learning approaches, we target an end-to-end training with direct feature learning, eliminating the pre-training stage in the common two-stage training scheme, and also examine how much of a performance boost image-level labels can further contribute in the supervised semantic segmentation task.

We highlight our main contributions in three-fold:

- We propose an end-to-end doubly (image- and pixel-levels) contrastive learning strategy for a lightweight semantic segmentation model to eliminate the pre-training stage in the conventional contrastive learning approach without requiring a large training batch size or a memory bank.
- Our training method achieves 1.34%p increase in mIoU measure from the baseline focal loss-only objective with the SwiftNet architecture (ResNet-18 backbone), running inference at up to 66.7 FPS in 2048×1024 resolution on a single Nvidia RTX 3080 Mobile GPU.
- We verify that replacing image-level supervision with self-supervision in our supervised contrastive objective achieves comparable performance when pre-trained with clear weather images.

2 Related Work

Supervised semantic segmentation. In categorizing each pixel into its representative semantic class, many previous studies have employed the encoder-decoder structure [9, 24, 25, 61]. An encoder-decoder network consists of an encoder module that gradually reduces the feature maps to capture enriched semantic features, and a decoder that also gradually proceeds to recover the spatial information. This structure allows for faster computation as it does not make use of dilated features in the encoder and can recover sharp object boundaries in the decoder [9, 61]. Extending this structure, DeepLabV3+ [7] adds a multi-scale aspect in the encoder as well as depth-wise separable convolution and atrous spatial pyramid pooling (ASPP) to obtain sharper segmentation outputs. While it lowers the computational cost overhead, it is still heavy in size (e.g., in floating point operations (FLOPs)) for practical deployment. On the other hand, SwiftNet [28] uses a real-time, efficient encoder that sequentially extracts and concatenates multi-scale pyramidal features, and ENet [29] is an even lighter model which extracts features in a single scale only.

Representation learning without supervision. The original idea of matching representations of a single image in agreement dates back to [6] and in the direction of preserving consistency in representations, [6, 22] advance the idea in semi-supervised settings. Previous attempts to adjust data representations appropriately to a specific task include learning with pretext tasks such as relative patch prediction [13], jigsaw puzzle solving [26], colorization [18] and rotation prediction [16]. A representative approach to learning generalized data representation by contrastive learning is SimCLR by Chen *et al.* [8], an augmentation-based contrastive learning method for consistent representations. It operates by augmenting the input with two random transformations and maximizing their mutual agreement, thereby achieving comparable performance to the supervised setting in image classification after pre-training on relevant pretext tasks. [9] extends SimCLR by further enhancing the representation learning capability with a larger model architecture, and validates its effectiveness when fine-tuned or distilled onto another smaller network.

Representation learning with supervision. SupContrast [20] extends SimCLR by directly incorporating image-level labels for the image classification task. In this supervised scheme, pairwise equivalence matching of representations based on their category labels enhances the push-pull effect in the feature space during training. In semantic segmentation, [40] introduces cross-image pixel-wise contrast to explicitly address intra-class compactness and inter-class dispersion phenomena to consider pixel-wise semantic correlations globally. However, as most contrastive representation learning methods are limited by a smaller mini-batch size than the number of necessary negative samples in practice, [40] stores all region (pooled) embeddings from training data in an external memory bank, transposing the loss from *pixel-to-pixel* to *pixel-to-region*. Another work in a semi-supervised setting [2] also uses a memory bank to store class-wise features learned from a teacher network, from which a subset of selected features are used as pseudo-labels for the student network in its pixel-wise contrast. Although both are trainable end-to-end owing to a memory bank, they do require considerable memory resources for storing feature embeddings.

Mixed supervision. Many works have sought to reduce human costs for acquiring labels by leveraging weak labels as well. [43] proposes a unified algorithm that learns from various forms of weak supervision including image-level tags, bounding boxes and partial labels to produce pixel-wise labels, and [10, 30] only utilize image-level labels as weak priors to generate accurate segmentation labels. As such, weak labels such as image-level labels have become useful in learning precise mappings from RGB to segmentation maps.

3 Methodology

3.1 Preliminaries

Self-Supervised Contrast. For a set of N randomly sampled image-label pairs, $\{\mathbf{x}_k, \mathbf{y}_k\}_{k=1\dots N}$, we prepare a corresponding *multi-viewed* set of augmented samples originating from the same sources, $\{\tilde{\mathbf{x}}_k, \tilde{\mathbf{y}}_k\}_{k=1\dots 2N}$, where each consecutive k^{th} - and $(k+1)^{\text{th}}$ -index sample pair originates from the same source. We take a batch of data containing two sets of N arbitrarily augmented samples as input to the self-supervised contrastive loss, following [8, 20] which stem from InfoNCE [17, 59]. We let $i \in I \equiv \{1, \dots, 2N\}$ be the anchor index, and the corresponding other augmented sample as $j(i)$ (*positive*). For $\mathbf{z}_i = \text{Proj}(\text{Enc}(\tilde{\mathbf{x}}_i)) \in \mathbb{R}^{D_{\text{proj}}}$, we apply the inner (dot) product and softmax operations on the i -th and the corresponding $j(i)$ -th projected encoded representations, as follows:

$$\mathcal{L}_{self} = \sum_{i \in I} \mathcal{L}_{self}^{(i)} = - \sum_{i \in I} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_{j(i)} / \tau)}{\sum_{a \in A(i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_a / \tau)} \quad (1)$$

where $A(i) \equiv I \setminus i$ and the remaining set of images $\{k \in A(i) \setminus j(i)\}$ contains $2(N-1)$ number of *negative* samples. $\tau \in \mathbb{R}^+$ is the temperature parameter for the softmax, whereby the higher the temperature value, the softer (lower) the logits output of the softmax function. For each anchor i , we use 1 positive pair and $2N-2$ negative pairs, thus in total of $2N-1$ terms to obtain the loss value.

Supervised Contrast. For *image-level* supervised contrast, there are present more than one sample belonging to each image class label. We thus take the average of all values over all *positives* for the anchor sample i , such that $P(i) \equiv \{p \in A(i) | \tilde{\mathbf{y}}_p = \tilde{\mathbf{y}}_i\}$. Note that the key difference between self-supervised and supervised contrastive approaches is that for each anchor, multiple *positives* and *negatives* from the same and different classes, respectively, are considered rather than different data augmentations of the same anchor. Following the loss objective in [20] which takes the summation outside the log operator, we define the image-level supervised contrastive loss as follows:

$$\mathcal{L}_{image} = \sum_{i \in I} \mathcal{L}_{image}^{(i)} = - \sum_{i \in I} \frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_p / \tau)}{\sum_{a \in A(i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_a / \tau)} \quad (2)$$

where $|P(i)|$ denotes the cardinality of the set of *positives*. Note that easy positives and negatives (*i.e.*, the dot product ~ 1) contribute to the gradient relatively smaller than the hard positives and negatives (*i.e.*, the dot product ~ 0). We refer to [20] for detailed proofs.

Pixel-level supervised contrastive loss follows a similar form as the image-level, except that the representation is on the each pixel rather than the entire sample, and *positives* and *negatives* come from the same image rather than from two randomly augmented images. As in Wang *et al.* [44], the pixel-wise contrast loss addresses two limitations in using the cross-entropy (CE) loss, in which (1) it penalizes predictions neglecting the pixel-wise semantic correlations, and (2) the relative relations among pixel-wise logits fail to directly supervise on the features. We formulate the pixel-level supervised contrastive loss as follows:

$$\mathcal{L}_{pixel} = - \frac{1}{|P(i)|} \sum_{i_p \in P(i)} \log \frac{\exp(\mathbf{i} \cdot \mathbf{i}_p / \tau)}{\sum_{i_a \in A(i)} \exp(\mathbf{i} \cdot \mathbf{i}_a / \tau)} \quad \forall i \in \{1, \dots, H \times W\}, \quad (3)$$

where $A(i) \equiv Q \setminus i$, in which Q represents the set of all pixels in the multi-viewed data batch and i denotes the pixel anchor index in a given image. $P(i)$ denotes pixel embedding collections of the *positive* samples for each pixel i in an image of size $H \times W$.

3.2 Doubly Contrastive Supervised Segmentation

We remark that our approach as shown in Fig. 1 is, in fact, also applicable to other supervised semantic segmentation tasks in which we have available ground truth pixel-level semantic labels as well as the corresponding image-level labels for those pixels in each image. Adverse weather conditions is one such case where we can readily acquire and exploit image-level contexts to promote contextual contrastive learning under unfavorable outdoor weather conditions.

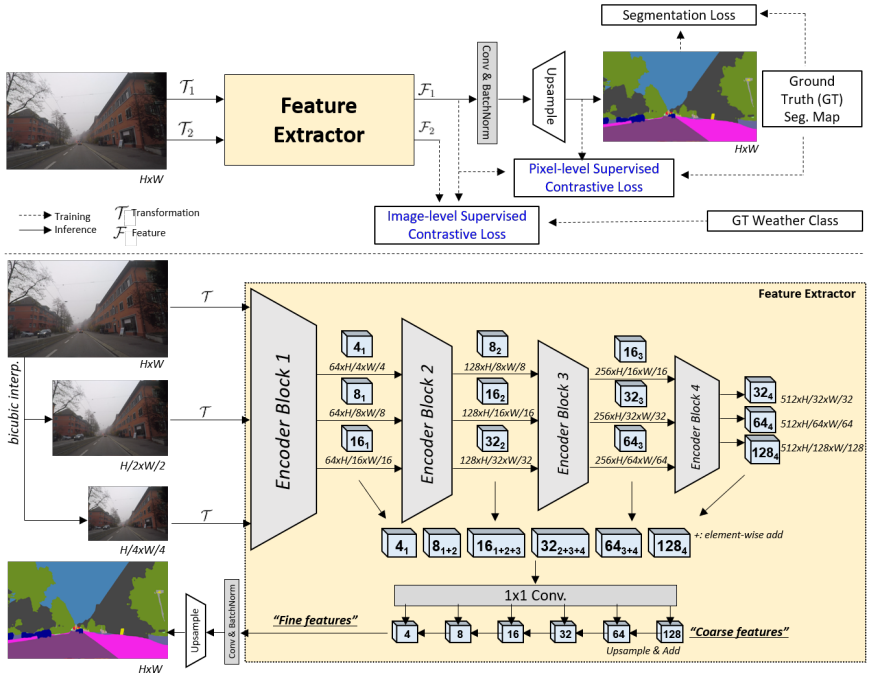


Figure 1: Overview of the training pipeline (top) and SwiftNet architecture (bottom).

The base model for our experiments is SwiftNet [28] with the ResNet-18 backbone to balance the trade-off between the accuracy and the scale of model parameters and computations. SwiftNet is advantageous for its light weight and effectiveness in self-driving scenarios that often require real-time inference speed as well as comparably high performance to those of large-scale models [66]. Moreover, its multi-scale pyramidal features extracted from the sequential encoder blocks provide scale-invariant features that capture spatial information more precisely than single-scale models [21].

Loss objective. Standard semantic segmentation networks have long been trained according to the pixel-wise CE loss given class-specific probability distributions by weighting all pixels equally and considering the frequency of pixels belonging to each class [28, 66]. This approach, however, is prone to overfitting and biased towards more frequent classes. In addition, the CE loss suffers from failing to recognize edges in pixel-level granularity. To remedy these problems, [23, 28] have employed class balancing by adopting a scaling coefficient term to weight each class in a less biased and balanced manner.

We thus follow [66] in order to weight the cross-entropy and the focal loss [23] with the class-balancing term as described in Eq. 4, where we take the inverse logarithm of the ratio of the frequency of each pixel appearance, $freq_c$, over the entire pixels in the dataset, N_p . We choose $\varepsilon = 1 \times 10^{-1}$ for numerical stability. The Euclidean distance transform (EDT) [45] serves to weight each pixel by the distances to the boundaries of other objects. That is, for each pixel p in a given image grid $G_{H \times W}$, we compute the L_2 distance to the nearest unmasked pixel q , *i.e.*, the closest valid pixel belonging to the set of classes, C as in Eq. 4.b. With both EDT and class-balancing terms in effect, the pixels of smaller objects and object boundaries are weighted greater than more frequently seen pixels such as those of *sky*

or *road* classes during training.

$$\mathcal{L}_{seg}(\phi(p), \hat{\phi}(p)) = -\delta(p)e^{\gamma(1-P_t)}\log(P_t), \quad (4)$$

with

$$\delta(p) = \underbrace{\log\left(1 + \varepsilon + \frac{freq_c(p)}{N_p}\right)^{-1}}_{\text{Class balancing}} \cdot \underbrace{\exp\left(-\frac{d_{EDT}(p)}{2\sigma_{EDT}}\right)}_{\text{EDT}}, \quad (4.a)$$

$$d_{EDT}(p) = \sum_C \min_{q \in C} \|p - q\|_2 \quad \forall p \in G_{H \times W}, \quad (4.b)$$

where P_t denotes the softmax probability for each semantic class, and $\phi(p)$ and $\hat{\phi}(p)$ are the ground truth and the predicted segmentation maps, respectively. We then combine the segmentation and the contrastive loss terms as the final loss objective as follows:

$$\mathcal{L} = \lambda_c \cdot (\mathcal{L}_{image} + \mathcal{L}_{pixel}) + \lambda_s \cdot \mathcal{L}_{seg}, \quad (5)$$

where $\lambda_c = 1/\mathcal{B}$, $\lambda_s = 1.2$ and \mathcal{B} is the batch size. For \mathcal{L}_{image} and \mathcal{L}_{pixel} , we further add stability during training by adding L_2 normalization to the inner product (*i.e.*, $\|\mathbf{z}_i \cdot \mathbf{z}_p\|_2$) to directly mimic the cosine similarity. The scaling coefficients, λ_c and λ_s , are experimentally determined so as to reduce the unfavorable effect of feature equivalence matching and place a greater weight upon the segmentation objective. We also scale the contrastive loss value empirically to be close to the segmentation loss value. In this combined objective, the image-level contrast complements the pixel-level contrast such that the semantic correlations among intra-image pixels of each image are consistent across images of different weather classes globally without a memory bank. For the following experiments that use self-supervision in place of image-level contrast, we simply replace \mathcal{L}_{image} with \mathcal{L}_{self} .

4 Experiments

4.1 Datasets and Evaluation Metric

We evaluate the proposed method using the training and validation sets of an urban street scene dataset, the Adverse Conditions Dataset with Correspondences for Semantic Driving Scene Understanding (ACDC) [35], with and without pre-training on another driving scene benchmark dataset, the Cityscapes [10] as described in Table 1.

Cityscapes is a road scene benchmark dataset containing stereo RGB images taken from the egocentric point of view of the driver under clear, daytime weather, with pixel-level annotations for 19 semantic classes of various roadside objects including bus and fence. ACDC is a smaller dataset labelled according to the Cityscapes categories, yet in adverse conditions (fog, nighttime, rain and snow). While other Cityscapes-like datasets [10, 12, 33, 34, 37, 44, 45] contain few to several hundred images biased toward specific adverse conditions, ACDC offers evenly-balanced number of images and the corresponding labels.

We evaluate the semantic segmentation performance by the mean Intersection-over-Union (IoU), also known as the Jaccard Index. IoU is often computed using the confusion matrix, from which we obtain true positive (TP), false positive (FP) and false negative (FN) scores for the predictions on the validation set, *i.e.*, $IoU = TP / (TP + FP + FN)$. The average IoU score computed over all fine category labels is denoted by mIoU, neglecting the *void* label.

Table 1: Dataset Description

Dataset	Input Modality	Resolution	Anno. (# Classes)	Weather Condition	Train	Val	Test
Cityscapes [10]	Stereo RGB	2048×1024	Fine (19)	Clear/Daytime	2,975	500	1,525
				All	1,600	406	2,000
ACDC [5]	Monocular RGB	1920×1080	Fine (19)	Fog	400	100	500
				Nighttime	400	106	500
				Rain	400	100	500
				Snow	400	100	500

4.2 Implementation Details

We performed our experiments on a single Nvidia RTX 3090 GPU with PyTorch 1.7.1, CUDA 11.1, and CUDNN 8.3.2. In order to account for scale variations, we augmented the data with the following transformation: random square crop (768×768) followed by scaling by a random value sampled from a uniform distribution $\sim U(0.5, 2.0)$. We did not apply transformations that could potentially harm the image quality in adverse weather conditions such as color jitter. We set the validation image width and height to those of *Cityscapes* (2048×1024) for fair comparison of the segmentation results. We pre-trained each model with ImageNet [52] and set $\gamma = 0.5$ in the focal loss, the temperature $\tau = 0.07$, and the feature dimension to $D_{proj} = 128$ for the supervised contrast. We trained the network on the ACDC dataset for 400 epochs with a batch size of 8, using an Adam [27] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.99$. We set the initial learning rate and the weight decay to 4×10^{-4} and 1×10^{-4} , respectively, and used a cosine annealing scheduler to decay the learning rate to 1×10^{-6} in the last epoch.

4.3 Evaluation

Quantitative results. We summarize the semantic segmentation results in Table 2, where there are 1.34%p and 1.33%p increases in mIoU from the baseline for the ResNet-18 and -34 backbones, respectively, for our method. Our proposed method generally outperforms in mIoU under each of the four adverse conditions, relative to the other experimented single and double contrasts. Our method is particularly effective for the *harder* conditions, namely *nighttime*, *rain* and *snow*, where we exploit the semantic correlations in the embeddings across the pixels of different conditions owing to the doubly contrastive objective. We remark that the difficulties in predictions in nighttime and rain stem from a significant number of darker pixels obstructing clear view of objects in the driving area and rain droplets distorting the camera view and the texture of objects at far distances, respectively. The performance under *fog* is not relatively lower than in clear weather due to low fog density. Further, we highlight that our experiments are accompanied with a batch size of 8 under a more practical training scenario. We expect higher performance gain with a batch size of 2048 as suggested in common contrastive learning works [24].

With pre-training on clear weather images, there are 0.86%p and 1.80%p improvement in mIoU for the ResNet-18 and -34 backbones for our method, respectively, after fine-tuning on the adverse weather images. Note that the experiment (f) achieves as high overall mIoU as the experiment (g), as well as in each weather-specific mIoU score. While our method has minor improvements for all conditions except fog with the ResNet-18, we observed noticeable increases with the ResNet-18 backbones.

Table 2: Semantic segmentation performance by adverse weather conditions using SwiftNet. Bb denotes backbone. The best results in **boldface** and the second best in underline.

Bb	Exp.	Loss (*: single contrast, †: double contrasts)	mIoU (%)	Fog	Nighttime	Rain	Snow
ResNet-18	(a)	Cross Entropy	62.63	68.95	45.71	63.32	65.43
	(b)	Focal (baseline)	64.04	71.59	47.84	63.90	65.57
	(c)	+Pixel-level Supervised Contrast only*	63.79	69.76	47.26	63.05	68.34
	(d)	+Self-supervised Contrast only*	63.91	<u>71.83</u>	48.02	62.24	<u>68.35</u>
	(e)	+Image-level Supervised Contrast only*	63.62	69.66	47.65	63.28	<u>67.10</u>
	(f)	+Self-supervised and Pixel-level Supervised Contrasts†	<u>65.07</u>	72.45	48.57	<u>63.95</u>	68.31
	(g)	+Image- and Pixel-level Supervised Contrasts† (<i>Ours</i>)	65.38	67.94	<u>48.56</u>	65.38	68.64
ResNet-34	(a)	Cross Entropy	65.02	73.85	47.11	65.20	66.94
	(b)	Focal (baseline)	67.00	74.44	49.38	<u>67.29</u>	70.64
	(c)	+Pixel-level Supervised Contrast only*	66.60	72.49	49.84	65.55	69.06
	(d)	+Self-supervised Contrast only*	<u>68.09</u>	76.99	50.14	66.45	69.56
	(e)	+Image-level Supervised Contrast only*	68.07	75.12	<u>50.76</u>	66.78	69.33
	(f)	+Self-supervised and Pixel-level Supervised Contrasts†	67.02	72.06	50.12	65.86	<u>70.83</u>
	(g)	+Image- and Pixel-level Supervised Contrasts† (<i>Ours</i>)	68.33	<u>75.19</u>	51.21	67.37	71.19
Bb	Exp.	Loss (w/ pre-training with clear weather)	mIoU (%)	Fog	Nighttime	Rain	Snow
ResNet-18	-	Focal (<i>Cityscapes</i>)	73.16	N/A (Clear weather only)			
	(a)	Cross Entropy	64.40	71.88	47.09	65.59	66.66
	(b)	Focal (baseline)	65.49	73.43	47.67	64.77	68.35
	(f)	+Self-supervised and Pixel-level Supervised Contrasts†	66.97	<u>74.05</u>	49.10	67.85	69.69
	(g)	+Image- and Pixel-level Supervised Contrasts (<i>Ours</i>)†	<u>66.24</u>	75.38	<u>48.82</u>	<u>65.79</u>	<u>68.42</u>
ResNet-34	-	Focal (<i>Cityscapes</i>)	73.80	N/A (Clear weather only)			
	(a)	Cross Entropy	68.38	75.99	49.70	67.92	71.49
	(b)	Focal (baseline)	69.46	76.94	50.28	69.78	71.69
	(f)	+Self-supervised and Pixel-level Supervised Contrasts†	<u>70.06</u>	76.11	<u>50.47</u>	70.68	72.89
	(g)	+Image- and Pixel-level Supervised Contrasts (<i>Ours</i>)†	70.13	<u>76.31</u>	53.59	<u>70.50</u>	<u>71.89</u>

Qualitative results. As shown in Fig. 2, our doubly contrastive method corrects false positive predictions from the focal loss (baseline) case, and allows to predict objects of a relatively large size more consistently. For instance, it removes noisy predictions on *road*, removes the fuzzy boundary between *road* and *sidewalk*, and fills in mis-classified pixels with the correct label for somewhat noisily predicted objects. Specifically in *rain*, it recovers a *pole* that seemed nonexistent in the baseline and corrects the pixels mis-classified as *sidewalk* to *road*. The effect of our method is more apparent when there is a region with a majority of pixels belonging to a single class, yet there remains mis-classified or noisy pixels in part.

Feature visualization. For an in-depth analysis of the learned features, we provide the class- and pixel-wise t-SNE visualizations for the baseline and our method in Fig. 3. For the class-wise features categorized into the four weather conditions, our doubly contrastive method better discriminates images of the *fog* and *night* classes from images of the other two. Images of these two categories are more clustered and densely located by themselves.

For pixel-wise features, we opted to visualize for the first batch of pixel-wise feature embeddings due to a memory constraint, where we observed that the features are more densely populated, thereby appearing darker in each respective color. This insinuates pixels belonging to the same class are more compactly located in the feature space. While grouping patterns seem more apparent for the baseline, our combination of image- and pixel-level supervised contrastive objectives finds a good balance in the feature space that yields higher semantic segmentation scores.

Ablation studies. We examine the effect of different feature extractors as well as the *fine* vs. *coarse* feature granularity as noted in Fig. 1. Table 3 presents the results and the corresponding computational costs when the SwiftNet feature extractor is replaced by that of




Figure 2: Two samples of semantic segmentation prediction results for each of the four weather conditions: fog, night, rain and snow (1 \rightarrow 4). *Please zoom in to see details.*

ENet and DeepLabV3+ (ResNet-50 backbone). While our method is effective with SwiftNet, it is not as effective as with ENet or with DeepLabV3+. We infer that this is due to the single-scale encoder for ENet, and the ASPP module for DeepLabV3+. While SwiftNet draws multi-resolution features sequentially from the previous scale features, DeepLabV3+ applies convolutions with different kernel sizes and pools altogether in one step. We applied our image-level contrast to the ENet features before upsampling and DeepLabV3+ features after ASPP followed by 1×1 convolution, respectively. In terms of memory, ENet is the smallest in size, yet it requires heavy memory space and is thus relatively slower. DeepLabV3+, on the other hand, is large in size even with ResNet-50 instead of its proposed Xception-71. Replacing the *fine* features with the *coarse* in SwiftNet yields a drop of 3.75% in mIoU for our method since *fine* features contain the fused output from the features from multi-resolution encoder blocks that provide geometrically richer contexts than the *coarse*.

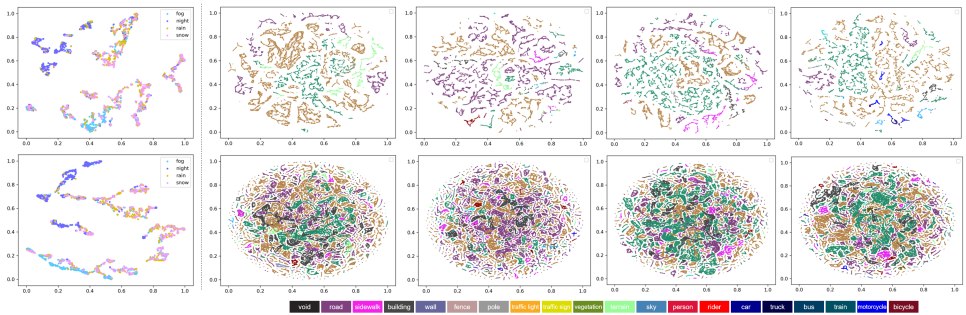


Figure 3: t-SNE visualizations for the trained features using SwiftNet (ResNet-18). Top row: focal loss only (baseline). Bottom row: Ours. Columns from left to right: weather class-wise features, and pixel-wise features for fog, nighttime, rain and snow, in order. *Void* class is not shown. *Please zoom in to see details.*

Table 3: Ablation study: semantic segmentation performance with different models and *coarse* features (2048×1024 resolution). *Coarse* features are marked with †; otherwise *fine*.

Model (Encoder type)	Exp.	mIoU (%)	Fog	Nighttime	Rain	Snow	GFLOPs	Params (M)	RTX 3080 Mobile Time (ms)	FPS
ENet [24] (Single-scale)	(a)	45.22	48.10	36.31	44.46	47.53	1.40	0.35	31	32.3
	(b)	50.45	55.14	38.70	49.44	53.44				
	(f)	50.78	55.91	38.96	50.88	52.39				
	(g)	49.32	53.64	37.85	49.86	51.17				
	(g)†	60.62	65.84	45.35	61.20	64.79				
SwiftNet (ResNet-18) [25] (Multi-scale pyramidal)	(a)	62.63	68.95	45.71	63.32	65.43	8.04	12.04	15	66.7
	(b)	64.04	71.59	47.84	63.90	65.57				
	(f)	<u>65.07</u>	72.45	48.57	63.95	<u>68.31</u>				
	(g)	65.38	67.94	<u>48.56</u>	65.38	68.64				
	(g)†	61.66	67.86	46.07	62.10	64.84				
SwiftNet (ResNet-34) [25] (Multi-scale pyramidal)	(a)	65.02	73.85	47.11	65.20	66.94	14.40	22.15	26	38.5
	(b)	67.00	<u>74.44</u>	49.38	<u>67.29</u>	70.64				
	(f)	<u>67.02</u>	72.06	<u>50.12</u>	65.86	<u>70.83</u>				
	(g)	68.33	75.19	51.21	67.37	71.19				
	(g)†	69.22	<u>74.54</u>	51.20	70.70	71.83				
DeepLabV3+ (ResNet-50) [1] (ASPP)	(a)	<u>69.69</u>	73.60	<u>51.45</u>	<u>72.18</u>	<u>72.29</u>	29.88	39.76	14	71.4
	(b)	70.07	75.18	52.49	73.22	71.50				
	(f)	69.22	73.95	50.54	70.42	73.24				
	(g)	69.04	<u>74.54</u>	51.20	70.70	71.83				
	(g)†	69.04	<u>74.54</u>	51.20	70.70	71.83				

5 Conclusion

We proposed an end-to-end doubly contrastive learning approach to semantic segmentation for self-driving under adverse weather. Our doubly contrastive method exploits image-level labels to semantically correlate RGB images taken under various weather conditions and pixel-level labels to obtain more semantically meaningful representations. In our method, the two supervised contrasts complement each other to effectively improve the performance of a lightweight model, without a need for pre-training or a memory bank to associate images across various weather conditions for global consistency. We hope this sheds further light on contrastive learning approaches for real-time deployment of self-driving systems.

Acknowledgments. This work was supported by the Institute for Information & Communications Technology Promotion (IITP) under Grant 2020-0-00440 through the Korean Government (Ministry of Science and ICT (MSIT); Development of Artificial Intelligence Technology that Continuously Improves Itself as the Situation Changes in the Real World).

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4981–4990, 2018.
- [2] Inigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8219–8228, 2021.
- [3] Jean Andrey and Sam Yagar. A temporal analysis of rain-related crash risk. *Accident Analysis & Prevention*, 25(4):465–472, 1993. ISSN 0001-4575. doi: [https://doi.org/10.1016/0001-4575\(93\)90076-9](https://doi.org/10.1016/0001-4575(93)90076-9). URL <https://www.sciencedirect.com/science/article/pii/0001457593900769>.
- [4] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.
- [5] Suzanna Becker and Geoffrey E Hinton. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161–163, 1992.
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [7] 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 (ECCV)*, pages 801–818, 2018.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [9] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255, 2020.
- [10] 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.
- [11] Dengxin Dai and Luc Van Gool. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *IEEE International Conference on Intelligent Transportation Systems*, 2018.
- [12] Dengxin Dai, Christos Sakaridis, Simon Hecker, and Luc Van Gool. Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *International Journal of Computer Vision*, 128(5):1182–1204, 2020.

- [13] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2051–2060, 2017.
- [14] Federal Motor Carrier Safety Administration. Code of federal regulations. Accessed May 23, 2022 [Online]. URL https://www.ecfr.gov/cgi-bin/retrieveECFR?gp=1&ty=HTML&h=L&mc=true&=PART&n=pt49.5.395#se49.5.395_12.
- [15] Pedro F Felzenszwalb and Daniel P Huttenlocher. Distance transforms of sampled functions. *Theory of computing*, 8(1):415–428, 2012.
- [16] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [17] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010.
- [18] 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.
- [19] Abbas Khan, Hyongsuk Kim, and Leon Chua. Pmed-net: Pyramid based multi-scale encoder-decoder network for medical image segmentation. *IEEE Access*, 9:55988–55998, 2021. doi: 10.1109/ACCESS.2021.3071754.
- [20] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- [21] Seung-Wook Kim, Hyong-Keun Kook, Jee-Young Sun, Mun-Cheon Kang, and Sung-Jea Ko. Parallel feature pyramid network for object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–250, 2018.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [24] 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.
- [25] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015.

- [26] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016.
- [27] World Health Organization. Road traffic injuries, 2021. URL <https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries>.
- [28] Marin Oršić and Siniša Šegvić. Efficient semantic segmentation with pyramidal fusion. *Pattern Recognition*, 110:107611, 2021.
- [29] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *CoRR*, abs/1606.02147, 2016. URL <http://arxiv.org/abs/1606.02147>.
- [30] Pedro O Pinheiro and Ronan Collobert. From image-level to pixel-level labeling with convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1713–1721, 2015.
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [32] 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 (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- [33] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, Sep 2018. URL <https://doi.org/10.1007/s11263-018-1072-8>.
- [34] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [35] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021.
- [36] Taek-Jin Song, Jongoh Jeong, and Jong-Hwan Kim. End-to-end real-time obstacle detection network for safe self-driving via multi-task learning. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–12, 2022. doi: 10.1109/TITS.2022.3149789.
- [37] Frederick Tung, Jianhui Chen, Lili Meng, and James J Little. The raincouver scene parsing benchmark for self-driving in adverse weather and at night. *IEEE Robotics and Automation Letters*, 2(4):2188–2193, 2017.
- [38] Abhinav Valada, Johan Vertens, Ankit Dhall, and Wolfram Burgard. Adapnet: Adaptive semantic segmentation in adverse environmental conditions. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4644–4651. IEEE, 2017.

- [39] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv e-prints*, pages arXiv-1807, 2018.
- [40] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020.
- [41] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7303–7313, 2021.
- [42] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268, 2020.
- [43] Jia Xu, Alexander G Schwing, and Raquel Urtasun. Learning to segment under various forms of weak supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3781–3790, 2015.
- [44] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020.
- [45] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steininger, and Gustavo Fernandez Dominguez. Wilddash-creating hazard-aware benchmarks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 402–416, 2018.
- [46] Pingping Zhang, Wei Liu, Hongyu Wang, Yinjie Lei, and Huchuan Lu. Deep gated attention networks for large-scale street-level scene segmentation. *Pattern Recognition*, 88:702–714, 2019.
- [47] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. Icnet for real-time semantic segmentation on high-resolution images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 405–420, 2018.