

Adaptive Sensitivity Analysis for Robust Augmentation against Natural Corruptions in Image Segmentation

Laura Zheng¹ Wenjie Wei¹ Tony Wu¹ Jacob Clements¹
Shreelekh Revankar^{1,2} Andre Harrison³ Yu Shen^{1,4} Ming C. Lin¹

Abstract

Achieving robustness in image segmentation models is challenging due to the fine-grained nature of pixel-level classification. These models, which are crucial for many real-time perception applications, particularly struggle when faced with natural corruptions in the wild for autonomous systems. While sensitivity analysis can help us understand how input variables influence model outputs, its application to natural and uncontrollable corruptions in training data is computationally expensive. In this work, we present an adaptive, sensitivity-guided augmentation method to enhance robustness against natural corruptions. Our sensitivity analysis on average runs **10x** faster and requires about **200x** less storage than previous sensitivity analysis, enabling practical, on-the-fly estimation during training for a model-free augmentation policy. With minimal fine-tuning, our sensitivity-guided augmentation method achieves improved robustness on both real-world and synthetic datasets compared to state-of-the-art data augmentation techniques in image segmentation. Code implementation for this work can be found at: <https://github.com/laurayuzheng/SensAug>.

1. Introduction

Segmentation models are crucial in many applications, but they often face unpredictable and uncontrollable natural variations that can degrade their performance. For instance, mobile applications using segmentation for image reconstruction may encounter diverse noises due to varying envi-

¹Department of Computer Science, University of Maryland, College Park, MD, USA ²Department of Computer Science, Cornell University, Ithaca, New York, USA ³DEVCOM Army Research Laboratory, Adelphi, Maryland, USA ⁴Adobe Research San Jose, California, USA. Correspondence to: Laura Zheng <lyzheng@umd.edu>.

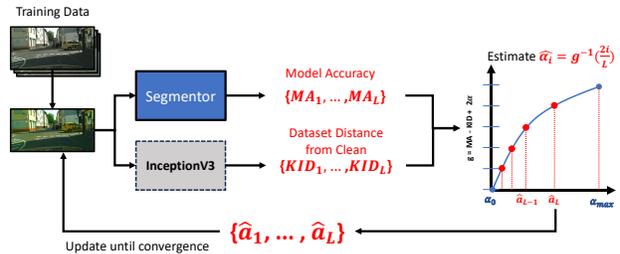


Figure 1. Overview of online sensitivity analysis. We conduct adaptive sensitivity analysis using our fast analysis algorithm after a warmup period on clean data, then solve for L discrete perturbation intensities, or levels, per perturbation type which the model is sensitive to. Finally, we augment training by sampling from the computed perturbation levels. Sampling weights are determined based off model performance on sensitive levels, where worse-performing levels are given higher probability of being sampled.

ronmental lighting, camera quality, and user handling. Similarly, autonomous vehicles and outdoor robots will need to operate under a wide range of adverse weather conditions that are difficult to simulate accurately. Even in medical imaging, where conditions are more controlled, factors such as slight movements can introduce blur, affecting segmentation results. While poor-quality examples can sometimes be discarded and re-captured, such solutions are costly or impractical, especially in large-scale, ubiquitous use cases, with limited resources, or during real-time inference (e.g., failure in a navigating robot). Addressing these natural corruptions is challenging because they are hard to capture ahead of time in a predictable and controllable way, simulate or parameterize, yet they significantly impact model performance.

One common approach to enhance robustness against such corruptions is data augmentation, which artificially increases the diversity of training data by applying transformations to existing samples. While data augmentation is convenient and resource-efficient, its effectiveness depends on selecting the most beneficial augmentations. Ideally, we should determine which augmentation a given model is most sensitive to and focus on those to improve performance—in other words, sensitivity analysis. However, traditional sensitivity analysis methods are computationally expensive

and resource-intensive (Shen et al., 2021), as shown in Table 1, making them impractical for large-scale or real-time applications. Existing methods like AutoAugment (Cubuk et al., 2019) and DeepAutoAugment attempt to optimize augmentation policies by training separate models, which adds significant overhead. Other state-of-the-art techniques rely on random augmentations (Cubuk et al., 2020; Muller & Hutter, 2021; Hendrycks et al., 2020), which are scalable but may not target the most impactful transformations.

In this paper, we propose a scalable, fast sensitivity-guided augmentation approach for robustifying segmentation models against natural corruptions, including those not explicitly involved during training. Our approach performs a lightweight, online sensitivity analysis during training to identify the geometric and photometric perturbations, shown to be effective as “basis perturbations” (Shen et al., 2021), to which the model is most sensitive. In contrast to (Shen et al., 2021), our sensitivity analysis is adaptive and significantly less resource intensive, allowing for practical implementation without the need for offline models or extensive computation. Figure 1 shows a high-level overview of our augmentation pipeline. Our method bridges the gap between the efficiency of random augmentation techniques and the effectiveness of policy-based augmentations guided by sensitivity analysis. Despite our focus on segmentation, our approach is general and can be applied to other tasks, architectures, or domains without significant modifications.

In experiments, we achieve up to a 6.20% relative mIoU improvement in snowy weather and up to a 3.85% relative mIoU improvement in rainy weather compared to the next-best method in zero-shot adverse weather evaluation on state-of-the-art architectures. We also show improvements on synthetic benchmarks and increased data efficiency (up to 200x) compared to other augmentation methods as the size of the training set changes.

Our contributions are summarized as follows:

1. An efficient *adaptive* sensitivity analysis method for *online model evaluation* that iteratively approximates model sensitivity curves for significant speedup;
2. A comprehensive, sensitivity-guided framework that systematically improves the robustness of learning-based segmentation models;
3. Evaluation and analysis of our method on *unseen* synthetically perturbed samples, *naturally corrupted* samples, and ablated contributing factors to robustification.

2. Related Works

Robustification Against Natural Corruptions. The effect of natural corruptions on deep learning tasks is a well-

explored problem, especially in image classification. Currently, image classification has a robust suite of benchmarks, including evaluation on both synthetic and natural corruptions (Hendrycks et al., 2020; Yi et al., 2021; Dong et al., 2020). Many works study correlations between image corruptions and various factors (Mintun et al., 2021; Hendrycks & Gimpel, 2017). Additionally, a popular approach to increasing robustness in the general case is through targeted adversarial training (Xiaogang Xu & Jia, 2021; Shu et al., 2021). Several approaches target model architecture (Schneider et al., 2020; Saikia et al., 2021; Myronenko & Hatamizadeh, 2020). Other approaches achieve robustness to natural corruptions via the data pipeline. Data augmentations are a popular method for increasing out-of-distribution robustness and many have now become standard practice (Geirhos et al., 2019; Rusak et al., 2020). Hendrycks et al. highlight that existing methods for generalization may not be consistently effective, emphasizing the need for robustness through addressing multiple distribution shifts (Hendrycks et al., 2021). In our work, we focus on studying and improving robustness in the context of semantic segmentation models due to natural corruptions using insights from previous work. Among findings from other works, we distinguish that our work focuses on improving natural corruption robustness in semantic segmentation, a common computer vision task with unique properties.

Data Augmentation Techniques. Data augmentation methods generate variants of the original training data to improve model generalization capabilities. These variants do not change the inherent semantic meaning of the image, and transformed images are typically still recognizable by humans. Within data augmentation methods, CutMix and AugMix are widely-used augmentation techniques that augment images by mixing variants of the same image (Hendrycks et al., 2020; Yun et al., 2019). Conversely, (Franchi et al., 2021) introduces segmentation-specific augmentation approaches which utilize superpixels, or clusters of similar pixels, to maintain semantic object information. Other data augmentation methods have utilized augmentation policies based on neural networks to select productive augmentations (Olsson et al., 2021; Cubuk et al., 2019; Zheng et al., 2022), while other works have explored data augmentation for domain-specific tasks (Zhao et al., 2019; Zhang et al., 2023). For example, (Zhao et al., 2019) explores learned data augmentation for biomedical segmentation tasks via labeling of synthesized samples with a single brain atlas. (Zhang et al., 2023) explores data augmentation in specifically brain segmentation via combining multiple brain scan samples, similarly to Augmix and Cutmix. However, this work is reliant on additional annotations to augment regions of interest. In our work, we present a generalizable augmentation technique and show that performance boosts generalize well out-of-the-box on several domains.

Method	SA Time	Data Gen Time	Storage
AdvSteer	90.0±15.5 min	~ 48 hours	2.4 TB
Ours	9.6±0.2 min	-	12 GB

Table 1. Runtime and Storage Comparison on Sensitivity Analysis of AdvSteer (Shen et al., 2021) vs. Ours. Our approach enables the practical use of sensitivity analysis in online training as an augmentation policy. We compute each mean and standard deviation value in “SA Time” across 4 runs. Each sensitivity analysis iteration computes curves for 24 different augmentations at 5 levels each, for a total of 120 evaluation passes. Advsteer requires an offline data generation stage for each dataset, whilst ours is entirely online. Computed SA Time does not factor in data generation time. Ours runs about $9.3\times$ faster and takes $200\times$ less storage in isolation.

3. Methodology

In general, sensitivity analysis examines how small fluctuations in the inputs affects the outputs of a system. In our augmentation approach, *the key idea is that sensitivity analysis can be used to sample augmentations uniformly with respect to “impact on model performance”, as opposed to sampling uniformly across the “parameterized augmentation space”.*

To quantify this for a given deep learning model, we need a metric for model performance and a metric for image degradation which is consistent across augmentation types. Choosing a model performance metric is straightforward; any bounded measure of accuracy (MA) where higher values are better suffices. As for the image degradation metric, we use Kernel Inception Distance (KID), introduced by (Bińkowski et al., 2018) to reduce bias towards sample size. At a high level, we use KID to measure the “distance” between an original dataset and its perturbed version. KID does so by passing both datasets through a generalized Inception model, and computing the square Maximum Mean Discrepancy (MMD) between their respective features. The reduced sample size bias of KID allows us to approximate the image degradation metric without iterating through the full validation set.

By sampling augmentations to which the model is sensitive, we can improve robustness. We define the sensitivity of the model to changes in augmentation intensity as the ratio of the change in model accuracy to the change in KID:

$$\text{sensitivity} = \frac{\Delta MA}{\Delta KID} \quad (1)$$

Our goal is to identify augmentation intensities that result in high sensitivity—that is, small changes in the augmentation (as measured by KID) lead to large changes in model performance (MA). This indicates that the model is particularly sensitive to those augmentations, and training on them could improve robustness. To formalize this, we seek to

Algorithm 1 Training with Sensitivity-Informed Augmentation.

Data: Training dataset X_t , Validation dataset X_v , Validation Rate r_v , SA Rate r_{SA}

Result: Trained semantic segmentation model

```

1  $N_V \leftarrow 0$ ; // Number of validation rounds
2  $f(\cdot) \leftarrow Identity(\cdot)$ ; // Aug transform
3 Initialize network weights  $\theta$ 
4 for  $i \leftarrow 1 \dots max\_iter$ ; // Training loop
5 do
6    $x_{ti} \leftarrow DataLoader(X_t)$  if  $p_f$  is initialized then
7      $f \sim p_f$ ; // Sample aug PDF
8   end
9    $x_{ti}^{aug} \leftarrow f(x_{ti})$ 
10  if  $i \% r_v == 0$  then
11    if  $i \% r_{SA} == 0$ ; // Update SA Curve
12    then
13      levels  $\leftarrow []$ ; // Store all  $\alpha$  values
14      metrics  $\leftarrow []$ ; // Store all metrics
15      for each augmentation type  $f$  do
16         $\alpha_f, acc_f \leftarrow SensitivityAnalysis(f, \theta)$ ;
17        // Appendix: Algorithm 2
18        levels.append( $\alpha_f$ ) metrics.append( $acc_f$ )
19      end
20      levels = levels.sort(); // Sort based on
21      // descending metrics
22       $p_f \leftarrow BetaBinom(idx(f), 0.75, 1.0)$ ;
23      // Categorical PDF by Acc
24    end
25    for  $x_{vi} \leftarrow DataLoader(X_v)$ ; // Validation
26    do
27      | Compute clean validation metrics
28    end
29  end

```

find a set of increasing, nontrivial augmentation intensities $\alpha_1 < \alpha_2 < \dots < \alpha_L$ that maximize sensitivity. We define the local changes in accuracy and KID between consecutive intensities as:

$$\Delta \widehat{MA}(\alpha_i, \alpha_{i-1}) = MA(\alpha_{i-1}) - MA(\alpha_i) \quad (2)$$

$$\Delta \widehat{KID}(\alpha_i, \alpha_{i-1}) = \frac{D_{KID}(x_{\alpha_i} \| x_{clean}) - D_{KID}(x_{\alpha_{i-1}} \| x_{clean})}{D_{KID}(x_{\alpha_{max}} \| x_{clean})} \quad (3)$$

Here, $MA(\alpha)$ is the model accuracy at augmentation intensity α , and $D_{KID}(x_{\alpha} \| x_{clean})$ is the KID between the augmented data at intensity α and the original clean data. The normalization in $\Delta \widehat{KID}$ ensures that KID values are comparable across different augmentation types.

We then formulate an objective function Q to find the set

of intensities that maximizes sensitivity while ensuring adequate spacing between them:

$$Q = \arg \max_{\alpha_1, \dots, \alpha_L} \min_{2 \leq i \leq L} \Delta \widehat{MA}(\alpha_i, \alpha_{i-1}) - \Delta \widehat{KID}(\alpha_i, \alpha_{i-1}) + \lambda(\alpha_i - \alpha_{i-1}) \quad (4)$$

In this equation, the term $\Delta \widehat{MA}(\alpha_i, \alpha_{i-1})$ represents the change in model accuracy between intensities α_{i-1} and α_i . We subtract $\Delta \widehat{KID}(\alpha_i, \alpha_{i-1})$ to favor intensity intervals where accuracy drops more than the image degradation increases, thus indicating higher sensitivity. Furthermore, the regularization term $\lambda(\alpha_i - \alpha_{i-1})$ (with $\lambda > 0$) encourages spacing between intensities, preventing them from being too close together. In our implementation, $\lambda = 2$.

Our objective seeks to maximize the minimum value of this expression across all intervals, ensuring that even the least favorable interval is optimized.

To simplify the optimization, we introduce a function $g(\alpha)$:

$$g(\alpha) = 1 - MA(\alpha) - \frac{D_{\text{KID}}(x_\alpha \| x_{\text{clean}})}{D_{\text{KID}}(x_{\alpha_{\max}} \| x_{\text{clean}})} + \lambda\alpha \quad (5)$$

The set of α values which fulfills Q has the following property: $g(\alpha_2) - g(\alpha_1) = g(\alpha_3) - g(\alpha_2) = \dots = g(\alpha_L) - g(\alpha_{L-1})$; in other words, optimal α values are produced at equal intervals along the function g . Since $g(\alpha)$ is approximately monotonically increasing (as $MA(\alpha)$ decreases and $D_{\text{KID}}(x_\alpha, x_{\text{clean}})$ increases with increasing α), and its values lie within a known range, we can approximate the solution as:

$$\alpha_i \approx g^{-1} \left(\frac{G_{\max} \cdot i}{L} \right), \quad i = 1, \dots, L \quad (6)$$

where G_{\max} is the maximum value of $g(\alpha)$ over the range of α , and g^{-1} is the inverse function. Since we choose $\lambda = 2$ in our implementation, $G_{\max} = 2$.

However, since we cannot explicitly compute g^{-1} due to $g(\alpha)$ being unknown in closed form, we iteratively estimate the values of α_i using methods like the Piecewise Cubic Hermite Interpolating Polynomial (PCHIP), which is a spline estimation technique. By sampling a few initial points and fitting an interpolating function, we can estimate the intensities that satisfy our objective. A proof for equal spacing can be found in Appendix Section A. We show the pseudocode for sensitivity analysis in Algorithm 2 of the appendix. Additionally, the iterative process for solving α

values is visualized in Appendix Figure 13. Below, we show the full training routine involving Sensitivity Analysis in Algorithm 1.

Resource improvement over previous work in sensitivity analysis. Previous sensitivity analysis methods (Shen et al., 2021) compute $g(\alpha)$ using a uniformly sampled set of α values across the entire augmentation space. This approach requires evaluating the model at many intensities and often necessitates offline generation of augmented datasets for each intensity and augmentation type. As a result, the storage complexity becomes the size of the original dataset multiplied by the number of augmentation types and intensities, leading to substantial storage demands.

In contrast, our method performs sensitivity analysis online during training and adaptively samples intensities based on the model’s responses. By estimating $g(\alpha)$ iteratively and focusing only on necessary intensities, we eliminate the need for pre-generating augmented datasets. As a result, our approach only adds about $0.2 * (\text{number of updates}) * (\text{evaluation time})$ amount of time to the total training pipeline, making the use of sensitivity analysis practical for on-the-fly augmentation policy during training.

4. Experiments

Hardware. Each experiment is conducted on four NVIDIA RTX A4000 GPUs and 16 AMD Epyc 16-core processors. Sensitivity analysis experiments are conducted on one GPU and 4 processors.

Experiment Setup. For evaluation on real-world corruptions and data efficiency, we train all experiments with the Segformer (Xie et al., 2021) backbone, a robust and state-of-the-art architecture for segmentation. Experiments in downstream fine-tuning from foundation model SAM (Kirillov et al., 2023), are restricted to their original ViT (Dosovitskiy et al., 2021) architectures as the backbone. All experiments, with the exception of downstream fine-tuning experiments, are trained for 160k iterations regardless of approach, and only the best-performing checkpoints by mIoU (mean Intersection-over-Union by class) are used for evaluation in results. Fine-tuning experiments initialized from foundation weights are trained for 80k iterations. Experiments within each table are run with the same hyperparameters with respect to learning configuration, initialization, and architecture. Additionally, nearly all models share the same set of augmentations, with the exception of IDBH (Li & Spratling, 2023), which uses an additional two augmentations (RandomFlip and RandomErase). We include these two additional augmentations in IDBH experiments to stay faithful to original release implementations by authors. We use official implementations for each method, and fix the random seed for each experiment such that they are repro-



Figure 2. **Special case on ACDC prediction: windshield wiper occlusion.** We observe a special case of natural corruptions in rainy weather which cannot be directly simulated by the existing set of perturbations: physical occlusion by windshield wipers. We compare our method to recent SOTA in augmentation for robustness, IDBH (Li & Spratling, 2023). *While IDBH involves random occlusion during training, ours does not. Our augmentation approach achieves comparable qualitative results with a smaller set of augmentations.* Additional visualizations can be found in Appendix Section F.

Method	Fog		Rain		Night		Snow	
	aAcc↑	mIoU↑	aAcc↑	mIoU↑	aAcc↑	mIoU↑	aAcc↑	mIoU↑
Baseline	89.70	55.10	87.41	42.82	54.39	14.89	83.23	41.22
AugMix	89.76	57.79	89.28	47.53	56.64	17.35	83.34	43.94
AutoAugment	77.06	56.18	75.52	42.66	57.14	20.65	71.83	40.94
RandAug	88.24	53.99	86.92	43.10	56.03	18.08	83.35	41.86
TrivialAug	85.79	55.16	84.35	41.26	54.52	17.02	77.99	42.64
VIPAug	92.04	60.28	89.19	46.41	61.10	17.40	85.72	45.04
IDBH	89.79	60.79	86.93	45.64	54.76	18.41	83.88	45.35
Ours	90.20	62.50	88.87	49.36	58.85	20.72	83.39	48.16

Table 2. **Evaluation of zero-shot adverse weather performance across data augmentation techniques.** We evaluate each data augmentation method across four different weather scenarios from the Adverse Conditions Dataset with Correspondences (ACDC) (Sakaridis et al., 2021) dataset. Each model is trained only with clean Cityscapes data with the Segformer (Xie et al., 2021) backbone. Our method, highlighted in grey, maintains the best performance across nearly all metrics for three out of four scenarios, with relative mIoU improvement over the next best method of up to 2.81% on fog, 3.85% on rain, and 6.20% on snow.

ducible. More hyperparameter details for experiments can be found in Appendix Section G. Full experiment configurations will be released alongside the code implementation for full reproducibility of results.

Metrics. We use three different metrics for evaluating the performance of a segmentation model: absolute pixel accuracy (aAcc), mean pixel accuracy (mAcc), and mean Intersection-over-Union (mIoU). Mean values are taken over object classes—thus, aAcc will be more susceptible to class imbalances, although it is the most intuitive.

4.1. Evaluation on Real-World Corruptions

To evaluate the robustness of our model in visual and graphics applications, we test on real-world adverse samples. While real-world adverse samples in most datasets are difficult to obtain, there are numerous real-world datasets for driving representing different cities and adverse weather scenarios.

We compare our results to seven methods: a baseline model where no augmentation is performed, AugMix (Hendrycks

et al., 2020), AutoAugment (Cubuk et al., 2019), RandAugment (Cubuk et al., 2020), and TrivialAugment (Muller & Hutter, 2021), VIPAug (Lee et al., 2024), and IDBH (Li & Spratling, 2023). The Inception model used to compute KID in our method is pre-trained on ImageNet; likewise, for policy-based augmentation techniques such as Augmix and AutoAugment, the policies are also based off of ImageNet. Different from other policy-based methods, our approach estimates the current model’s sensitivity to perturbations relative to the Inception model pre-trained on ImageNet and utilizes the information in augmentation sampling; no additional “policy” parameters are trained. On real-world dataset evaluation for unseen weather and domain gap scenarios, our method shows improvements over the next best performing model across almost all metrics. We include a qualitative visualization of our model versus several other methods in Figure 7 of the appendix, which shows inference on a rainy weather sample. Amongst all methods, a common failure mode is the presence of windshield wipers in rainy weather. A visualization of this can be found in Appendix Section J.

A break-down the performance on the ACDC dataset by

Dataset	Type	Method	Clean		Basis Aug		AdvSteer		IN-C	
			aAcc \uparrow	mIoU \uparrow						
ADE20K	General	TrivialAug	75.420	32.580	69.559	27.083	41.783	9.188	61.495	18.668
		IDBH	76.220	33.950	72.752	30.651	40.557	9.475	61.971	19.091
		Ours	76.110	33.790	74.285	31.922	43.075	9.628	61.280	18.721
VOC2012	General	TrivialAug	90.090	57.900	87.837	52.340	75.350	20.338	82.884	36.080
		IDBH	90.610	60.570	89.262	56.876	69.843	20.810	81.819	36.933
		Ours	90.800	61.140	89.555	58.183	69.690	21.470	82.519	38.834
POTSDAM	Aerial	TrivialAug	84.360	67.820	77.649	55.763	55.817	34.282	55.866	36.967
		IDBH	84.280	68.690	79.392	63.757	22.675	14.975	46.413	30.123
		Ours	84.550	68.450	82.590	66.065	44.817	29.983	54.275	36.416
A2I2Haze	UGV	TrivialAug	98.730	69.180	97.317	51.800	85.598	22.225	97.363	46.502
		IDBH	98.680	69.300	98.346	64.615	85.545	19.490	97.368	45.970
		Ours	98.790	70.290	98.613	67.919	89.482	21.843	97.407	49.805
Cityscapes	Driving	TrivialAug	95.570	74.300	86.117	56.952	69.785	30.593	82.664	44.332
		IDBH	95.530	73.930	93.160	68.052	71.932	29.388	83.041	44.225
		Ours	95.780	75.530	94.305	71.539	68.468	28.070	82.435	45.066
Synapse	Medical	TrivialAug	98.890	62.000	97.939	49.237	97.243	32.182	98.425	51.512
		IDBH	99.150	67.720	98.912	63.504	95.143	29.760	98.486	53.475
		Ours	99.250	71.380	99.082	68.828	90.282	30.310	96.779	56.013

Table 3. Performance evaluation of our method vs. SOTA on synthetic scenarios across 6 different datasets. We evaluate our method and SOTA on ADE20K (Zhou et al., 2019), VOC2012 (Everingham et al., 2012), POTSDAM (for Photogrammetry & Sensing), A2I2Haze (Narayanan et al., 2023), Cityscapes (Cordts et al., 2016), and Synapse (Landman et al., 2015) datasets, across three synthetic corruption scenarios: individual basis augmentations (Basis Aug), compositions of photometric augmentations produced by sensitivity analysis in Adversarial Steering (AdvSteer) (Shen et al., 2021), and the synthetic augmentation benchmark ImageNet-C (IN-C) (Hendrycks & Dietterich, 2019). Our method consistently achieves improved performance on synthetic corruption benchmarks while still maintaining or even improving clean evaluation accuracy.

weather type in Table 2. In total, the ACDC dataset has four different weather scenarios: Fog, Rain, Night, and Snow, where the largest relative boost over next-best method, IDBH (Li & Spratling, 2023), (6.20%) is in Snow scenarios. In three out of four weather categories, our method outperforms other methods, with the exception of Night scenarios. AugMix achieves higher aAcc but lower mIoU than our method on Rain scenarios possibly due to class imbalances, such as the large number of pixels classified as “sky”. While the total # of correct pixels is higher on AugMix, our method outperforms when averaged by class, on mIoU. Night scenario visibility corruption stems from lack of lighting, as opposed to the other three, which may have more differences in object appearances and blurring effects. While our method does not perform worse in mIOU, we do perform worse in aACC. This may suggest that the failure mode of our method in Night scenarios are due to smaller objects covering less pixel space. Performance on both ACDC and the India Driving (IDD) (Varma et al., 2019) datasets across multiple methods can be found in Section C of the Appendix.

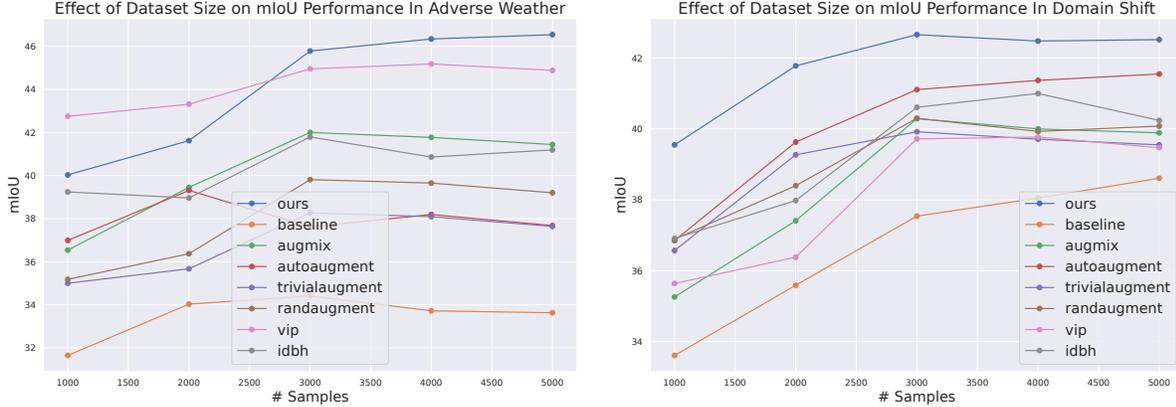
Special case: co-occurrence of windshield wipers and rainy weather. In the ACDC dataset, the rainy scenario evaluation set contains co-occurrences with windshield-wiper occlusion. This case is interesting in that occlusions are not included in any experiments except those of IDBH. In qualitative results, we observe that our method handles windshield wiper occlusions just as well, if not better, than IDBH. In Figure 2, we show an example of this, where our method shows comparatively less artifacts in the building and sky,

despite not having been trained on occlusion (RandomErase) augmentations.

4.2. Evaluation on Datasets

The results of previous experiments show the efficacy of our method in context of driving domains. In this experiment, we demonstrate that our method also shows improvements across several datasets and visual computing domains compared to SOTA.

We evaluated our method on six semantic segmentation datasets, ranging from generic everyday scenes to application-specific domains like driving or medicine. ADE20K (Zhou et al., 2019), VOC2012 (Everingham et al., 2012), POTSDAM (for Photogrammetry & Sensing), Cityscapes (Cordts et al., 2016), Synapse (Landman et al., 2015), and A2I2Haze (Narayanan et al., 2023). In Table 3, we show mIoU performance of our method versus the next-best augmentation technique, the SOTA baseline. We evaluate on clean data and three different synthetic scenarios: individual transformations from the basis augmentations at uniform parameter intervals (Basis Aug), the combined perturbation benchmark from (Shen et al., 2021) (AdvSteer), and ImageNet-C (IN-C) (Hendrycks & Dietterich, 2019) corruptions. *On the synthetic benchmark ImageNet-C (Hendrycks & Dietterich, 2019), our model achieves improved scores, particularly in the robotics and medical domains.* Our method performed worse primarily in the AdvSteer benchmark of Table 3, notably for Cityscapes and Synapse. This may be due to the sheer intensity of



(a) # Samples vs. ACDC mIoU performance.

(b) # Samples vs. IDD mIoU performance.

Figure 3. Comparison of Ours vs. Other SOTA Methods: Ours (top, blue) outperforms all others in performance as the number of samples increases, while other methods plateau on both (a) adverse weather data (ACDC) and domain shift (IDD).

Method	ViT+SAM		
	aAcc↑	mAcc↑	mIoU↑
Baseline	84.93	62.84	52.20
AugMix	84.69	63.25	54.18
AutoAugment	85.17	61.28	53.11
RandAug	85.16	59.33	51.95
TrivialAugment	84.87	59.92	50.58
VIPAug	84.29	61.00	51.93
IDBH	85.14	62.82	54.35
Ours	85.37	65.18	54.84

Table 4. Performance on ACDC when fine-tuning downstream segmentation with SAM. We show additional comparisons when initialized with SAM weights, similarly to results in Table 8.

benchmark corruption—the AdvSteer benchmark applies a combination of intense perturbations (not the same as the augmentations used during training), resulting in an extreme case from the original distribution. This may be related to degraded performance on Night scenarios in ACDC evaluation, as both scenarios heavily corrupt visibility based on color. Examples of the AdvSteer benchmark corruptions can be found in Appendix Section K. We emphasize that our method is not necessarily bound to image segmentation—we find similar boosts in performance in preliminary experiments with classification (see Appendix Section 12).

4.3. Downstream Finetuning with Foundation Models

A popular choice for boosting feature robustness is fine-tuning downstream tasks from foundation models. In these experiments, we examine how our approach can complement robustness provided by foundation models when fine-tuning on downstream tasks. We first initialize a distilled SAM (Kirillov et al., 2023) model on the ViT-Small (ViT-S) architecture, then fine-tune on the semantic segmentation

Method	aAcc↑	mIoU↑	mAcc↑
Baseline	94.19	66.67	75.28
AugMix	93.99	66.07	73.60
AutoAugment	93.84	64.99	72.09
RandAug	92.60	59.37	66.03
TrivialAugment	93.55	65.03	71.71
VIPAug	92.98	63.67	70.86
IDBH	93.62	65.49	74.96
Ours	93.88	68.03	75.96

Table 5. Domain adaptation; fine-tuning on a small target domain dataset. We show performance on the validation set of the ACDC Snow dataset, after training for 20k iterations on the ACDC Snow training set. All experiments are initialized with a Segformer-b0 model pre-trained on Cityscapes.

task with Cityscapes. We choose Cityscapes due to the availability of real-world corrupted images (ACDC and IDD) to evaluate on. In our experiments, we observe the highest performance on our method all three metrics. While the largest boost in robustness stem from robust foundation model features, our results suggest that our method can complement approaches centered around model architecture (such as Segformer). Additional results on downstream finetuning with DinoV2 can be found in Table 8 of the Appendix.

4.4. Domain Adaptation With Small Datasets

To evaluate the efficacy of our approach in domain adaptation settings, we finetune Segformer Cityscapes models on the training set of the ACDC snow data, then evaluate on the validation set of the ACDC snow data. The main difference from previous experiments here is that the model has been trained on target domain data, rather than being evaluated in a zero-shot setting. These results can be found in Table 5. Interestingly, the baseline approach seems to work second-best for domain adaptation, behind our method.

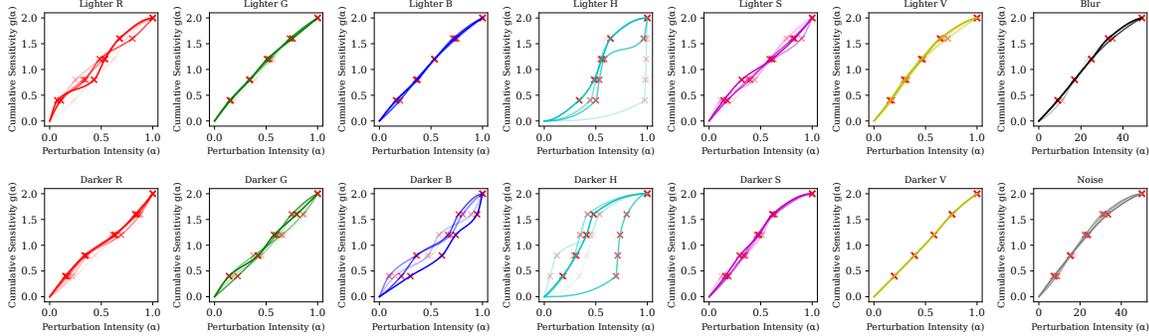


Figure 4. Sensitivity curves (g values) for color channels, noise, and blur throughout training of Cityscapes. We visualize how the estimated cumulative sensitivity curve, Equation 5, changes for RGB, HSV, Gaussian blur, and Gaussian noise during augmented training. In this plot, the most recent curve is opaque, while others decrease in opacity in order of recency. The red X markers indicate the values at which α values are selected (horizontal axes). Surprisingly, most curves remain largely stagnant throughout training, with the exception of Hue in HSV (teal, center), which changes drastically as the model generalizes. This may suggest that Hue is a highly adversarial factor in model generalization. Ablation study results in Table 6 support this, where the model trained without photometric augmentations demonstrate a significant decrease in performance.

Method	Clean		Basis Aug		AdvSteer		IN-C		ACDC	
	aAcc \uparrow	mIoU \uparrow								
Baseline	95.610	75.130	92.042	65.319	62.040	21.995	79.437	38.362	78.49	37.54
Ours \sim_g	95.780	75.500	93.405	68.877	71.070	27.997	83.032	44.385	78.13	43.69
Ours \sim_p	95.740	75.210	92.544	69.002	64.907	22.437	80.817	40.876	75.74	37.97
Ours \sim_{Warmup}	95.830	75.430	94.458	71.891	69.138	28.472	84.438	45.849	79.78	44.66
Ours $\sim_{Uniform}$	95.740	75.200	94.304	71.213	69.678	27.235	85.135	46.219	80.95	43.17
Ours	95.790	75.100	94.439	71.665	70.605	28.895	83.844	45.617	80.13	44.67

Table 6. Ablation study results comparing different variants of our method. We compare: (1) a baseline trained with no augmentations, (2) a variant of our method that only augments with photometric augmentations (Ours \sim_g), (3) a variant of our method that only uses geometric augmentations (Ours \sim_p), (4) a variant of our method trained without clean training warm-up, (5) a variant of our method with uniform augmentation (Ours $\sim_{Uniform}$) of computed sensitivity analysis values α , and (6) our full method combining informed probability sampling, and adaptive sensitivity analysis, and all augmentation types (Ours).

We suspect that the degradation in performance is due to overfitting; with smaller datasets and less challenging augmentations, the model may be more prone to overfitting than generalization. Our method may mitigate this by sampling evenly difficult augmentations with respect to the current model state.

4.5. Data Efficiency

We also analyze data efficiency of our method in comparison to other data augmentation methods by training various Segformer models with varying training dataset sizes. For each augmentation method, we train five models with training dataset sizes of 1000, 2000, 3000, 4000, and 5000 samples from the Cityscapes dataset. We plot the progression of mIoU (Minimum Intersection over Union) performance (higher the better) on (a) adverse weather data (ACDC) and (b) the domain shift setting (IDD), as shown in Figure 3. Our method, in blue, shows consistent improvement on adverse weather and domain shift evaluation with increasing number of samples, and maintains best mIoU performance across each # of samples slice, suggesting that *our method is more*

data efficient than others. Interestingly, not all methods show increased robustness to adverse weather as the number of samples increases for training, indicating that *in some cases, scaling data may not necessarily mean increased robustness*.

4.6. Ablation Study

We examine several variants of our method to determine the impact of individual components in an ablation study: a baseline trained only with random cropping, a variant of our method using only geometric augmentations, a variant of our method using only photometric augmentations, a variant of our method without clean training warm-up, and a variant of our method using uniform sampling instead of the Beta-Binomial sampling described in Algorithm 1. Uniform sampling of augmentation parameters computed with sensitivity analysis decreases generalization to both synthetic and real-world corruption benchmarks by small margins. In addition, training without clean warm-up produces similar results to that with warm-up, suggesting that warm-up is optional. In our case, warming up with clean evaluation

reduces the total number of sensitivity analysis updates, making warm-up with clean evaluation marginally less resource expensive (~ 0.5 GPU hours total). Interestingly, while clean performance remains largely the same across all models, the largest decrease in performance on unseen corruption benchmarks comes from the lack of photometric augmentations.

4.7. Color Channel Sensitivity Across Training

Sensitivity analysis can also be useful for interpretable analysis on-the-fly during training, aside from being used in augmentations directly. To analyze how sensitivity changes with respect to different channel values, we plot each g curve computed from Equation 5 for RGB, HSV, Noise, and Blur channels across training for our Cityscapes experiments in Figure 4. Note that the curves in this figure are based on individual color channels and are separate from those used during training for generalization analysis purposes. Computed α values will center towards regions with higher sensitivity relative to the current model. From this visualization, we observe that Hue curves (teal, center) are most volatile during training, with most sensitive augmentation parameters falling towards α values close to 1.0 in the beginning of training. As the model generalizes, the Hue curve converges slowly towards α values centered around 0.5, similarly to other curves. *This suggests that Hue may be a significant adversarial factor in model robustness compared to other channels.*

5. Discussion and Conclusion

In this paper, we present a method for sensitivity-informed augmented training for semantic segmentation. Our work is inspired by applications in real-time perception systems such as robotics, where natural corruptions may be abundant at inference time. Our method combines the information granularity of sensitivity analysis-based methods and the scalability of data augmentation methods, which run on-the-fly during training. In our results, we show that our method achieves improved robustness on zero-shot real-world adverse weather and domain shift scenarios, in addition to improvements on synthetic benchmarks like ImageNet-C. Additionally, evaluation on real world datasets show clear improvements over current SOTA methods for augmentation. Our model can complement other approaches for model robustness such as architecture design and downstream fine-tuning with foundation models.

Acknowledgement: This research is supported in part by the U.S. Army Research Labs Cooperative Agreement on “AI and Autonomy for Multi-Agent Systems”.

Impact Statement

The goal of this work is to introduce an augmentation framework which enhances generalization to naturally-occurring corruptions. Our work is driven by problems in real-time systems such as robotics. To the best of our knowledge, combating natural corruptions has only positive implications for the systems which it is a concern; for example, autonomous vehicles may be more robust to adverse weather or unexpected scenarios.

References

- Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Cubuk, E. D., Zoph, B., Mané, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation strategies from data. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 113–123, 2019. doi: 10.1109/CVPR.2019.00020.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. Randaugment: Practical automated data augmentation with a reduced search space. In *Advances in Neural Information Processing Systems*, pp. 18613–18624, 2020.
- Dong, Y., Fu, Q.-A., Yang, X., Pang, T., Su, H., Xiao, Z., and Zhu, J. Benchmarking adversarial robustness on image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houselby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>, 2012.
- for Photogrammetry, I. S. and Sensing, R. Potsdam: 2d semantic labeling contest. <https://www.isprs.org/education/benchmarks/UrbanSemLab/2d-sem-label-potsdam.aspx>.

- Franchi, G., Belkhir, N., Ha, M. L., Hu, Y., Bursuc, A., Blanz, V., and Yao, A. Robust semantic segmentation with superpixel-mix, 2021. URL <https://arxiv.org/abs/2108.00968>.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., and Brendel, W. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bygh9j09KX>.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Hkg4TI9xl>.
- Hendrycks, D., Mu, N., Cubuk, E. D., Zoph, B., Gilmer, J., and Lakshminarayanan, B. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*, 2020.
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., Song, D., Steinhardt, J., and Gilmer, J. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 8340–8349, October 2021.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4015–4026, 2023.
- Landman, B., Xu, Z., Iglesias, J. E., Styner, M., Langerak, T. R., and Klein, A. Multi-atlas labeling beyond the cranial vault - workshop and challenge. <https://www.synapse.org/#!Synapse:syn3193805/wiki/89480>, 2015.
- Lee, I., Lee, W., and Myung, H. Domain generalization with vital phase augmentation. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’24/IAAI’24/EAAI’24. AAAI Press, 2024. ISBN 978-1-57735-887-9. doi: 10.1609/aaai.v38i4.28070. URL <https://doi.org/10.1609/aaai.v38i4.28070>.
- Li, L. and Spratling, M. W. Data augmentation alone can improve adversarial training. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=y4uc4NtTWaq>.
- Mintun, E., Kirillov, A., and Xie, S. On interaction between augmentations and corruptions in natural corruption robustness. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 3571–3583. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/1d49780520898fe37f0cd6b41c5311bf-Paper.pdf.
- Muller, S. G. and Hutter, F. Trivialaugment: Tuning-free yet state-of-the-art data augmentation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 754–762, 2021.
- Myronenko, A. and Hatamizadeh, A. Robust semantic segmentation of brain tumor regions from 3d mris. In Crimi, A. and Bakas, S. (eds.), *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, pp. 82–89, Cham, 2020. Springer International Publishing. ISBN 978-3-030-46643-5.
- Narayanan, P., Hu, X., Wu, Z., Thielke, M. D., Rogers, J. G., Harrison, A. V., D’Agostino, J. A., Brown, J. D., Quang, L. P., Uplinger, J. R., Kwon, H., and Wang, Z. A multi-purpose realistic haze benchmark with quantifiable haze levels and ground truth. *IEEE Transactions on Image Processing*, 32:3481–3492, 2023. doi: 10.1109/TIP.2023.3245994.
- Olsson, V., Tranheden, W., Pinto, J., and Svensson, L. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1369–1378, January 2021.
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H. V., Szafraniec, M., Khalidov, V., Fernandez, P., HAZIZA, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.-Y., Li, S.-W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., and Bojanowski, P. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=a68SUt6zFt>.

- Rusak, E., Schott, L., Zimmermann, R. S., Bitterwolf, J., Bringmann, O., Bethge, M., and Brendel, W. A simple way to make neural networks robust against diverse image corruptions. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J.-M. (eds.), *Computer Vision – ECCV 2020*, pp. 53–69, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58580-8.
- Saikia, T., Schmid, C., and Brox, T. Improving robustness against common corruptions with frequency biased models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 10211–10220, October 2021.
- Sakaridis, C., Dai, D., and Van Gool, L. 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.
- Schneider, S., Rusak, E., Eck, L., Bringmann, O., Brendel, W., and Bethge, M. Improving robustness against common corruptions by covariate shift adaptation. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 11539–11551. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/85690f81aadcd1749175c187784afc9ee-Paper.pdf.
- Shen, Y., Zheng, L., Shu, M., Li, W., Goldstein, T., and Lin, M. Gradient-free adversarial training against image corruption for learning-based steering. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 26250–26263. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/dce8af15f064dlaccb98887a21029b08-Paper.pdf.
- Shu, M., Shen, Y., Lin, M. C., and Goldstein, T. Adversarial differentiable data augmentation for autonomous systems. In *International Conference on Robotics and Automation (ICRA)*, pp. 14069–14075, 2021.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Varma, G., Subramanian, A., Namboodiri, A., Chandraker, M., and Jawahar, C. Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1743–1751, 2019. doi: 10.1109/WACV.2019.00190.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. Caltech-ucsd birds-200-2011. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- Xiaogang Xu, H. Z. and Jia, J. Dynamic divide-and-conquer adversarial training for robust semantic segmentation. In *ICCV*, 2021.
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J. M., and Luo, P. Segformer: Simple and efficient design for semantic segmentation with transformers. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 12077–12090. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/64f1f27bf1b4ec22924fd0acb550c235-Paper.pdf.
- Yi, C., YANG, S., Li, H., peng Tan, Y., and Kot, A. Benchmarking the robustness of spatial-temporal models against corruptions. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=MQLMIrm3Hv5>.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *International Conference on Computer Vision (ICCV)*, 2019.
- Zhang, X., Liu, C., Ou, N., Zeng, X., Zhuo, Z., Duan, Y., Xiong, X., Yu, Y., Liu, Z., Liu, Y., and Ye, C. Carvemix: A simple data augmentation method for brain lesion segmentation. *NeuroImage*, 271:120041, 2023. ISSN 1053-8119. doi: <https://doi.org/10.1016/j.neuroimage.2023.120041>. URL <https://www.sciencedirect.com/science/article/pii/S1053811923001878>.
- Zhao, A., Balakrishnan, G., Durand, F., Guttag, J. V., and Dalca, A. V. Data augmentation using learned transformations for one-shot medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Zhao, H., Shi, J., Qi, X., Wang, X., and Jia, J. Pyramid scene parsing network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Zheng, Y., Zhang, Z., Yan, S., and Zhang, M. Deep autoaugmentation. In *ICLR*, 2022.

Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., and Torralba, A. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019.

Appendix

A. Proof on Equal Spacing

1. $\alpha_0 = 0$ represents an augmentation intensity of 0, i.e. a clean image.
2. $MA(0)$ approaches 1 as models get better, but to be more precise, the function $g(\alpha)$ should be

$$g(\alpha) = MA(0) - MA(\alpha) - \frac{D_{\text{KID}}(x_\alpha \| x_{\text{clean}})}{D_{\text{KID}}(x_{\alpha_{\text{max}}} \| x_{\text{clean}})} + \lambda\alpha.$$

3. Q is the set of points $\{\alpha_1, \dots, \alpha_L\}$ that maximizes the following:

$$\begin{aligned} g' &= \min_{1 \leq i \leq L} \widehat{MA}(\alpha_i, \alpha_{i-1}) - \Delta \widehat{KID}(\alpha_i, \alpha_{i-1}) + \lambda(\alpha_i - \alpha_{i-1}) \\ &= \min_{1 \leq i \leq L} g(\alpha_i) - g(\alpha_{i-1}). \end{aligned}$$

4. From the definition of $g(\alpha)$, we note

$$\begin{aligned} g(0) &= 0 \\ g(\alpha_L) &= G_{\text{max}} = (MA(0) - MA(\alpha_L)) - 1 + \lambda\alpha_L. \end{aligned}$$

In our implementation, we normalize MA so that $g(\alpha_L) = 2$.

Proof of equal-spacing:

We will prove that $g' = G_{\text{max}}/L$, i.e. given the set of α_i values that fulfills Q , the set of $g(\alpha_i)$ are equally spaced along the y -axis of the function $g(\alpha)$. BWOC, assume $g' > G_{\text{max}}/L$. Then, $\forall i, g(\alpha_i) - g(\alpha_{i-1}) \geq g' > G_{\text{max}}/L$. Taking the sum over i yields a contradiction:

$$G_{\text{max}} = g(\alpha_L) - g(0) = \sum_{i=1}^L g(\alpha_i) - g(\alpha_{i-1}) > L \left(\frac{G_{\text{max}}}{L} \right) = G_{\text{max}}.$$

Thus, $g' \leq G_{\text{max}}/L$. If we assume $g(\alpha)$ is continuous over $[0, \alpha_L]$, then by the Intermediate Value Theorem, $\forall i, \exists a_i$ such that

$$g(a_i) = \frac{G_{\text{max}} \cdot i}{L}.$$

Then, $\forall i, g(\alpha_i) - g(\alpha_{i-1}) = G_{\text{max}}/L$, so the maximum $g' = G_{\text{max}}/L$ can be attained for a specific set of points Q . In the paper, we assume $g(\alpha)$ is strictly monotonically increasing because $MA(\alpha)$ decreases and $D_{\text{KID}}(x_\alpha, x_{\text{clean}})$ increases as α increases. Based on this, we obtain formula 6.

B. Sensitivity Analysis Pseudocode

Algorithm 2 Fast Sensitivity Analysis

Data: Number of levels L , Uncertainty threshold ϵ

Result: Perturbation Levels $\{\alpha_1, \dots, \alpha_{L-1}\}$

```

28  $g(\alpha) \leftarrow$  Equation 5 points  $\leftarrow \{(0, 0), (\alpha_L, 2)\}$  loop
29    $\hat{c} \leftarrow$  PCHIP(points) for  $i \leftarrow 1 \dots L - 1$  do
30      $\alpha_i \leftarrow$  Estimate( $\hat{c}$ ,  $2i/L$ ) ( $y_l, y_u$ )  $\leftarrow$  Estimate upper and lower y-values of  $\hat{c}$  at  $x = \alpha_i$   $\hat{c}_l \leftarrow$ 
       PCHIP(points.insert( $y_l$ ))  $\hat{c}_u \leftarrow$  PCHIP(points.insert( $y_u$ ))  $\alpha_{i_l} \leftarrow$  Estimate( $\hat{c}$ ,  $y_l$ )  $\alpha_{i_u} \leftarrow$  Estimate( $\hat{c}$ ,  $y_u$ )  $\epsilon_i$ 
        $\leftarrow (\alpha_{i_u} - \alpha_{i_l})/2$ 
31   end
32    $\alpha^*, \epsilon^* \leftarrow$  Choose level with max  $\epsilon_i$  if  $\epsilon^* < \epsilon$  then Break loop;
33   points.insert( $(\alpha^*, g^*(\alpha^*))$ )
34 end;
    
```

C. Overall Performance on ACDC and IDD

Method	Weather // ACDC			Domain // IDD		
	aAcc \uparrow	mIoU \uparrow	mAcc \uparrow	aAcc \uparrow	mIoU \uparrow	mAcc \uparrow
Baseline	76.31	35.48	47.36	85.82	38.44	59.14
AugMix	79.57	40.90	52.74	86.52	40.50	62.43
AutoAugment	70.29	39.31	54.18	85.79	40.74	62.24
RandAug	78.46	39.07	52.32	85.54	38.99	59.82
TrivialAug	75.50	38.56	53.62	85.23	39.61	61.04
IDBH	78.65	41.67	53.65	86.49	40.48	61.74
VIPAug	81.85	42.26	53.59	85.94	38.97	58.79
Ours	80.16	45.45	57.58	85.76	40.33	63.03

Table 7. **Evaluation results on Unseen Real World Driving Datasets.** We conduct zero-shot evaluation of Cityscape models on both ACDC (Sakaridis et al., 2021) and IDD (Varma et al., 2019) datasets, which represent adverse weather and domain transfer to India respectively. Our method achieves clear improvements compared to other methods which require chained, more computationally expensive augmentations or external augmentation models in terms of generalization to real world scenarios, with relative mIoU improvement up to 9.07% on ACDC compared to the next-best, IDBH.

D. Downstream Finetuning with DinoV2

Method	ViT+DinoV2		
	aAcc \uparrow	mAcc \uparrow	mIoU \uparrow
Baseline	77.65	45.83	32.70
Augmix	79.99	51.63	41.38
AutoAugment	81.18	55.93	43.65
RandAugment	80.42	54.02	43.25
TrivialAugment	82.56	54.27	43.58
IDBH	84.45	60.22	48.69
Ours	84.13	62.92	49.82

Table 8. **Performance of Cityscapes models on unseen ACDC weather evaluation set across different augmentation methods, when fine-tuned from DinoV2 (Oquab et al., 2024) with ViT (Dosovitskiy et al., 2021) backbone.**

E. Downstream Finetuning with SAM

Method	Fog		Rain		Night		Snow	
	aAcc \uparrow	mIoU \uparrow						
Baseline	93.96	68.98	91.59	58.53	65.73	29.41	90.54	56.27
AugMix	94.38	70.49	92.18	60.66	62.47	25.35	90.46	58.60
AutoAugment	94.16	69.18	92.57	56.63	63.66	26.69	90.98	58.19
RandAug	93.92	67.67	91.70	55.30	65.19	24.00	90.47	56.70
TrivialAug	93.71	66.45	91.06	53.09	60.57	20.25	90.56	54.67
IDBH	94.12	70.61	92.06	55.71	65.25	28.22	90.98	59.67
VIPAug	93.92	67.54	91.06	55.49	63.00	25.01	89.88	57.85
Ours	94.34	70.98	92.66	57.92	66.93	27.24	91.08	57.32

Table 9. **Evaluation on ACDC adverse weather performance with SAM downstream finetuning.**

F. Qualitative Results on Window Wiper Occlusion in ACDC

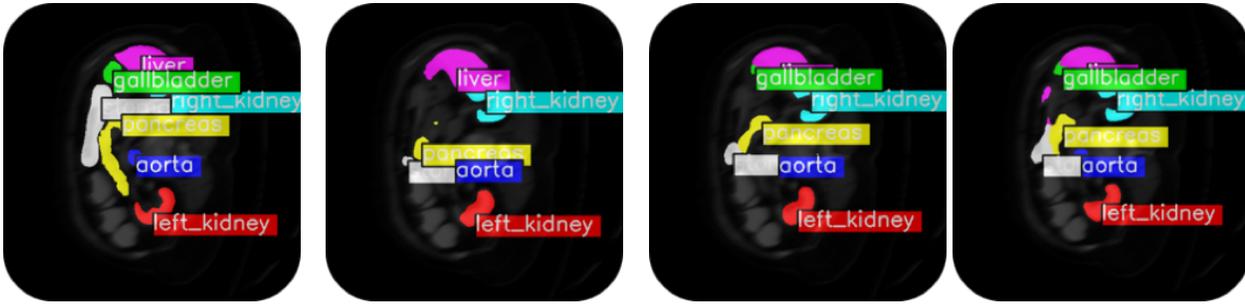
G. Detailed Experiment Hyperparameters

Let $AA = \{\text{RandomCrop, Contrast, Equalize, Posterize, Rotate, Solarize, Shear X, Shear Y, Translate X, Translate Y, Color, Contrast, Brightness, Sharpness}\}$ be the standard set of augmentations used with AutoAugment.

Method	Max Iters	LR	Optimizer	Augmentations	Batch Size	Backbone
Baseline	160,000	6e-05	AdamW	RandomCrop	1	SegFormer-b0
Augmix	160,000	6e-05	AdamW	AA	1	SegFormer-b0
AutoAugment	160,000	6e-05	AdamW	AA	1	SegFormer-b0
RandAug	160,000	6e-05	AdamW	AA	1	SegFormer-b0
TrivialAug	160,000	6e-05	AdamW	AA	1	SegFormer-b0
IDBH	160,000	6e-05	AdamW	AA \cup {RandomFlip, RandomErasing}	1	SegFormer-b0
Ours	160,000	6e-05	AdamW	AA	1	SegFormer-b0
$r_v = 1600;$ $r_{SA} = 9600;$						
Ours	160,000	6e-05	AdamW	AA	1	SegFormer-b0
$Warmup = 6400$						

Table 10. Experiment hyperparameters for Table 7 and Table 3. All experiments are trained under similar hyperparameter settings, with each evaluation conducted on the highest-performing mIoU checkpoint. In comparisons, we prioritize official implementations released by authors and avoid re-implementations. Additionally, most comparisons use the same set of augmentations to ours, with the exception of IDBH (Li & Spratling, 2023), whose original implementation includes RandomFlip and RandomErasing. For all experiments, we use the SegFormer-b0 backbone (Xie et al., 2021), which is a recent state-of-the-art segmentation-specialized architecture.

H. Qualitative Results on Synapse



(a) Ground truth. (b) TrivialAugment Prediction. (c) IDBH Prediction. (d) Our Prediction.

Figure 6. Qualitative evaluation on multi-organ segmentation with motion blur corruption. We show predictions on a motion-blurred sample from the Synapse (Landman et al., 2015) dataset for TrivialAugment (b), IDBH (c), and Our method (d), against the ground truth (a). Our method is able to segment right and left kidneys, liver, and aorta accurately. In contrast, the TrivialAugment prediction is unable to distinguish both kidneys.

I. Qualitative Results on Rainy Data



(a) AutoAugment Prediction. (b) IDBH Prediction. (c) Our Prediction.

Figure 7. Qualitative comparison on snowy urban driving sample between AutoAugment (Cubuk et al., 2020), IDBH (Li & Spratling, 2023), and Ours. In this example, each method (AutoAugment, IDBH, Ours) is trained on clean Cityscapes data representing sunny weather, then evaluated on adverse weather samples. Despite not having rainy data in the training set, our method is able to segment the driving noticeably clearer than other methods. In particular, other methods consistently struggle to segment the vehicle confidently.

J. Special Case: Windshield Wiper Occlusion

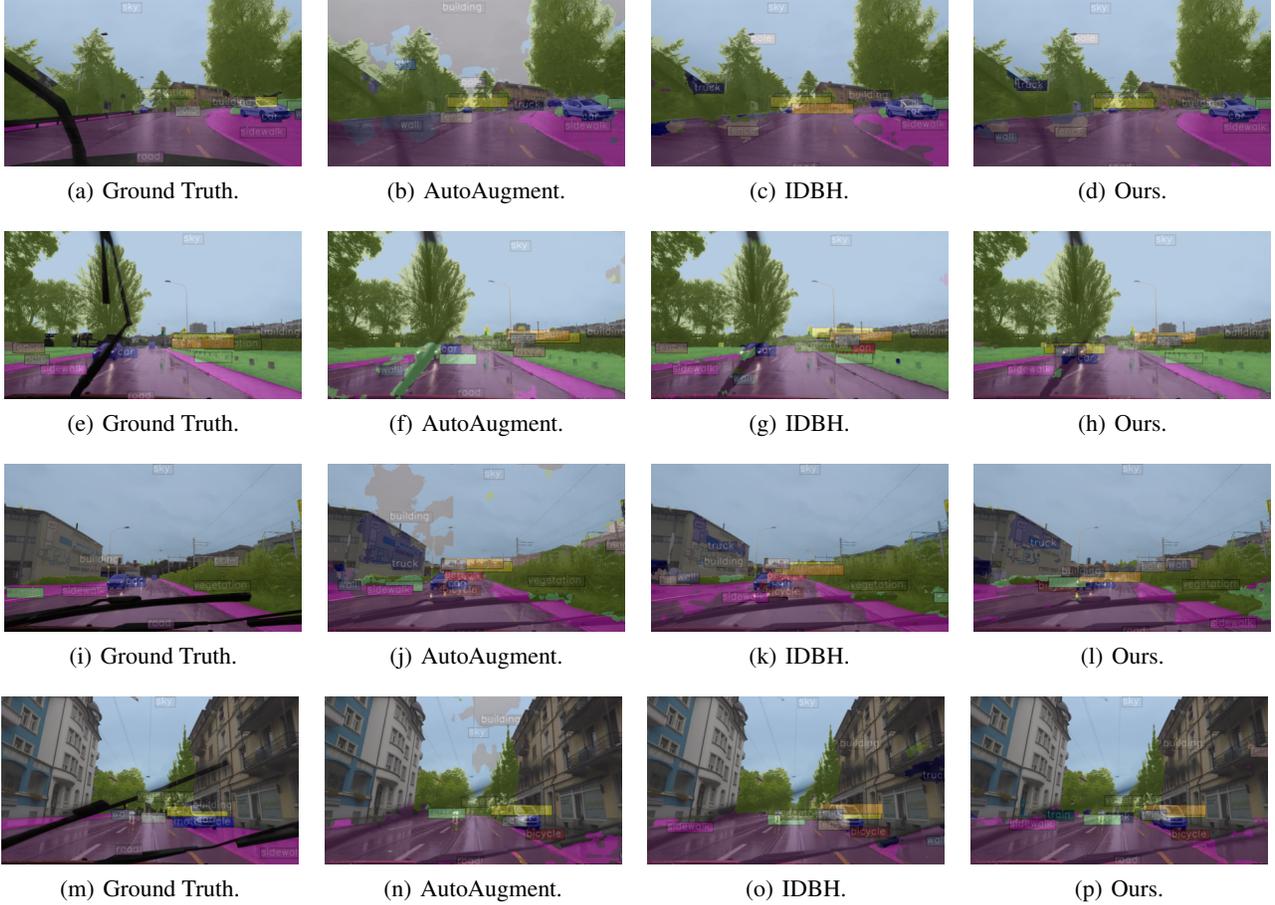


Figure 8. More examples of special case on ACDC prediction: windshield wiper occlusion.

J.1. Details on Basis Augmentations

Previous work in robustification showed that learning with a set of “basis perturbations” (BP) significantly improved zero-shot evaluation against unseen corruptions (Shen et al., 2021) for image classification and regression tasks, such as vehicle steering prediction. The intuition behind basis perturbations is that the composition of such perturbations spans a much larger space of perturbations than may be observed in natural corruptions; observed zero-shot performance boosts on unseen corruptions subsequently might be attributed to learning a model which is robust to basis perturbations. In our method, we extend this concept and introduce a more generalized and larger set of basis perturbations in sensitivity analysis to determine the most productive augmentation during training.

Let $D = \{Positive, Negative\}$ describe the set of augmentations applied in either a positive (lighter) direction or negative (darker) to either one channel of an image or a parameter of an affine transformation applied to an image.

Let $P = \{R, G, B, H, S, V\}$ describe the set of channels in RGB and HSV color spaces which may be perturbed; in other words, these augmentations are *photometric*.

Then, let $G = \{ShearX, ShearY, TranslateX, TranslateY, Rotate\}$ denote affine, or *geometric*, transformations which are parameterized by a magnitude value.

Finally, let $Z = \{Noise, Blur\}$ be the set of augmentations not applied along channel dimensions. Specifically, we use Gaussian Noise and Gaussian Blur.

Thus, the set of all basis augmentations A_B used in robustification is $A_B = \{D \times P + G + Z\}$.

To compute lighter or darker channel augmentations of RGB or HSV channels, we use linear scaling. Let the range of a channel be $[v_{\min}, v_{\max}]$. For lighter channel augmentations, we transform the channel values v_C by an intensity factor α like so:

$$v'_C = \alpha v_{\max} + (1 - \alpha) \cdot v_C$$

Likewise, for darker channel augmentations, the transformation can be described like so:

$$v'_C = \alpha v_{\min} + (1 - \alpha) \cdot v_C$$

The default values are $v_{\min} = 0$ and $v_{\max} = 255$. For H channel augmentations, we set the maximum channel values to be 180. For V channel augmentations, we set the minimum channel values to be 10 to exclude completely dark images.

Affine transformations can be represented as a 3×3 matrix, which, when multiplied with a 2-dimensional image, produces a geometrically distorted version of that image. Affine transformation matrices are typically structured in the form:

$$M = \begin{bmatrix} 1 & Shear_X & T_x \\ Shear_Y & 1 & T_y \\ 0 & 0 & 1 \end{bmatrix}$$

for shear and translation transformations. For rotations where the center of the image is fixed as the origin point $(0, 0)$, the transformation matrix is defined as:

$$M_{rot} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

To account for padded values in images after affine transformations, we zoom in images to the largest rectangle such that padded pixels are cropped out.

All augmentations are parameterized by a magnitude value ranging from 0 to 1. A magnitude value of 1 corresponds to the most severe augmentation value. More details on exact parameter value ranges can be found in the appendix. Conversely, a magnitude value of 0 produces no changes to the original image, and can be considered an identity function. We account for the symmetry of these augmentation transformations by considering both positive values and negative values as separate augmentations. The fast adaptive sensitivity analysis algorithm introduced in the next section relies on the property that increasing magnitude corresponds to increasing “distance” between images. Thus, augmentations cannot simply span the value ranges -1 to 1, and we separate them instead to different augmentations (positive and negative).

We apply these augmentations on-the-fly in online learning rather than generating samples offline. Doing so greatly reduces the offline storage requirement by one order of magnitude. Suppose L intensity levels are sampled for each basis augmentation. Then, offline generation of perturbed data requires up to $L \times 2 \times (|P| + |G|) + 2 = 24L$ additional copies of the original clean dataset. *With online generation, we avoid offline dataset generation entirely* and only need the original clean dataset to be stored, similar to standard vanilla learning.

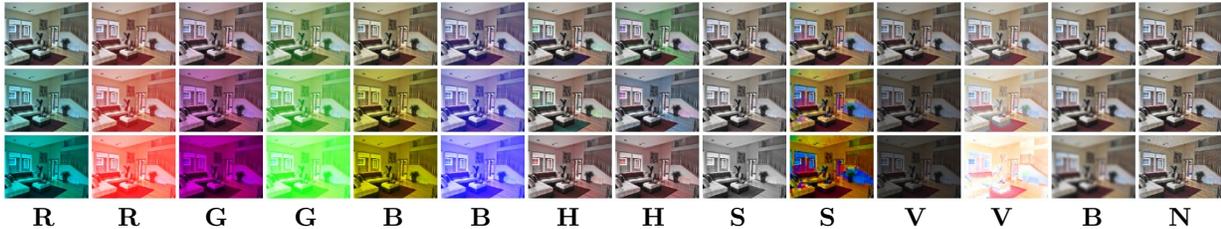


Figure 9. Visualization of each photometric augmentation transformation on a bedroom image. Up \uparrow indicates the “lighter”, positive direction and \downarrow indicates the “darker”, negative direction. “B” and “N” indicate blur and noise, respectively.



Figure 10. Visualization of various geometric augmentations applied to a sample image of a house. We use the following geometric transformations in our sensitivity analysis scheme, which are also analogous to the set of transformations used by other methods (Cubuk et al., 2019; Zheng et al., 2022). Up arrows indicate augmentation in the *positive*, or left, direction, while down arrows indicate augmentation in the *negative*, or right, direction.

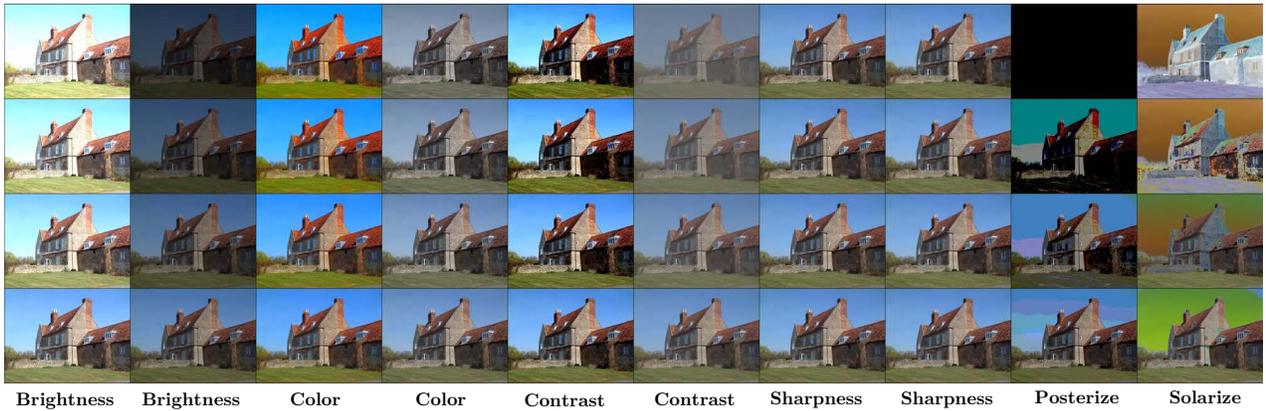


Figure 11. Additional augmentation types used in sensitivity analysis, which are used in other methods such as AutoAugment. While these photometric transformations are used in other methods, the transformations also overlap with the photometric transformations shown in Figure 9, namely HSV perturbations. However, we still conduct sensitivity analysis evaluation on these transformations for completion.

K. AdvSteer Benchmark Examples



Figure 12. AdvSteer benchmark examples.

L. Clean Performance on Different Backbones

Method	PSPNet (Zhao et al., 2017)			SegFormer (Xie et al., 2021)		
	aAcc \uparrow	mAcc \uparrow	mIoU \uparrow	aAcc \uparrow	mAcc \uparrow	mIoU \uparrow
Baseline	63.770	48.695	35.715	86.825	57.280	48.365
Augmix	94.770	74.400	66.740	95.520	81.430	73.390
AutoAugment	95.130	77.210	69.630	95.550	81.390	73.820
RandAugment	95.060	76.770	69.360	95.610	82.390	74.560
TrivialAugment	95.090	75.930	68.620	95.640 <td 83.210	75.130	
Ours	95.100	79.320	71.840	95.880	84.070	76.330

Table 11. Comparison of clean evaluation performance across different augmentation methods on Cityscapes. We evaluated our sensitivity-informed augmentation method against popular benchmarks on PSPNet and SegFormer. The baseline represents training with no augmentations.

M. Results on CUB Dataset for Classification

Method	InceptionV3			
	Clean	Basis Aug	AdvSteer	IN-C
Baseline	41.647	15.965	3.679	20.501
Augmix	35.865	15.274	4.810	20.394
AutoAugment	16.793	7.219	2.575	8.158
TrivialAugment	33.914	13.338	4.229	17.586
RandAugment	36.624	15.466	4.821	19.345
Ours	47.670	18.122	5.276	21.842

Table 12. Performance on CUB (Wah et al., 2011) dataset with InceptionV3 (Szegedy et al., 2016) backbone.

N. Fast Sensitivity Analysis Illustration

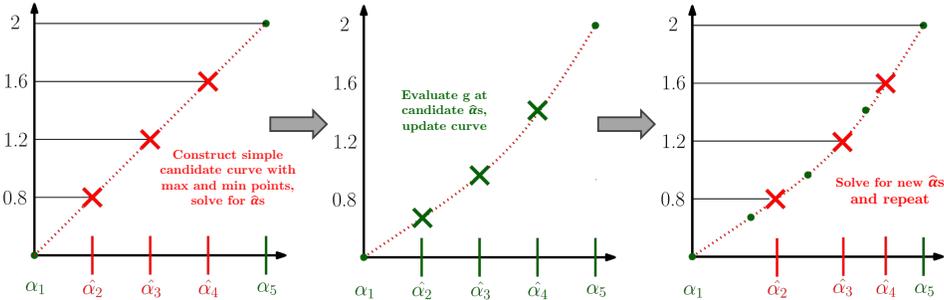


Figure 13. Illustration of fast sensitivity analysis. Each iteration of the fast sensitivity can be intuitively visualized. Since we can assume general monotonicity of the curve, we first initialize a candidate curve (a line in the first iteration). We solve for the candidate perturbation levels $\hat{\alpha}$ based on the solution in Equation 6. In the next step (middle), we evaluate the candidate level with the greatest uncertainty and adjust the candidate curve, the dotted red line, using PCHIP on the evaluated levels, which are guaranteed to be true points along the function g from Equation 5. In the next step (right), we use the new curve and solve for new candidate levels, repeating the process in the previous two steps until the maximum uncertainty of any candidate level values falls below a threshold of 0.05.

O. Sensitivity Analysis Computed Curve Comparison

Perturb	Method	p_1	p_2	p_3	p_4
R_{\uparrow}	Baseline	0.100	0.300	0.500	0.700
	Adaptive	0.149	0.253	0.399	0.604
G_{\uparrow}	Baseline	0.100	0.200	0.400	0.600
	Adaptive	0.103	0.204	0.395	0.619
B_{\uparrow}	Baseline	0.200	0.300	0.500	0.700
	Adaptive	0.146	0.328	0.551	0.788
R_{\downarrow}	Baseline	0.200	0.400	0.600	0.800
	Adaptive	0.225	0.503	0.625	0.803
G_{\downarrow}	Baseline	0.200	0.400	0.600	0.800
	Adaptive	0.256	0.447	0.607	0.812
B_{\downarrow}	Baseline	0.200	0.500	0.700	0.800
	Adaptive	0.231	0.450	0.594	0.730
H_{\uparrow}	Baseline	0.100	0.300	0.400	0.900
	Adaptive	0.268	0.406	0.508	0.809
S_{\uparrow}	Baseline	0.200	0.500	0.600	0.800
	Adaptive	0.243	0.439	0.589	0.744
V_{\uparrow}	Baseline	0.200	0.400	0.600	0.700
	Adaptive	0.193	0.360	0.517	0.680
H_{\downarrow}	Baseline	0.200	0.400	0.500	0.600
	Adaptive	0.279	0.433	0.548	0.699
S_{\downarrow}	Baseline	0.200	0.400	0.600	0.900
	Adaptive	0.199	0.344	0.562	0.847
V_{\downarrow}	Baseline	0.200	0.400	0.600	0.800
	Adaptive	0.197	0.397	0.594	0.797
<i>blur</i>	Baseline	9	19	25	35
	Adaptive	9	17	23	31
<i>noise</i>	Baseline	10	15	20	30
	Adaptive	6.4	12.4	17.7	26.9

Table 13. Comparison of computed perturbation levels using a baseline (Shen et al., 2021) sensitivity analysis method versus our adaptive method. p_5 is 1 for all RGB/HSV channels, 49 for blur, and 50 for noise. In previous work, each perturbation level is chosen from a certain number of sampled, discretized values. Additionally, these perturbed datasets are generated offline in an additional step before training. Our fast sensitivity analysis enables sensitivity analysis to be performed on the fly during training, and offers much more dynamic, accurate, and descriptive sensitivity curves.

O.1. KID vs. FID Relative Error Comparison with Scaling Sample Sizes

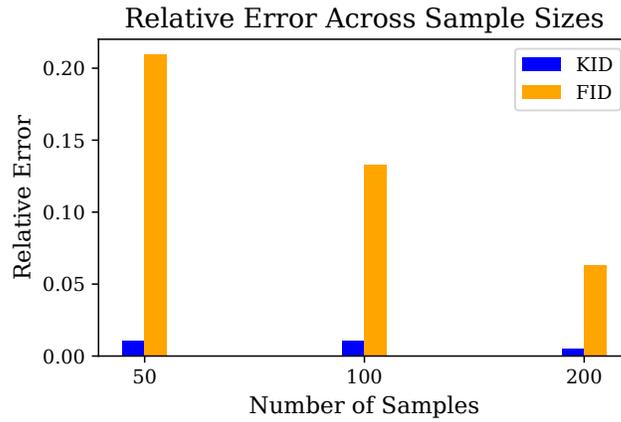


Figure 14. **Relative error of KID and FID over several sample sizes.** We plot the relative error of computed KID and FID values over several sample sizes, with the reference value being the computed value for each at 500 samples. From this, we can see that FID is significantly biased toward the number of samples used for evaluation. We can reduce the evaluation of KID values in sensitivity analysis by a notable fraction due to this property.

P. Train-time Evaluation on Perturbed Datasets

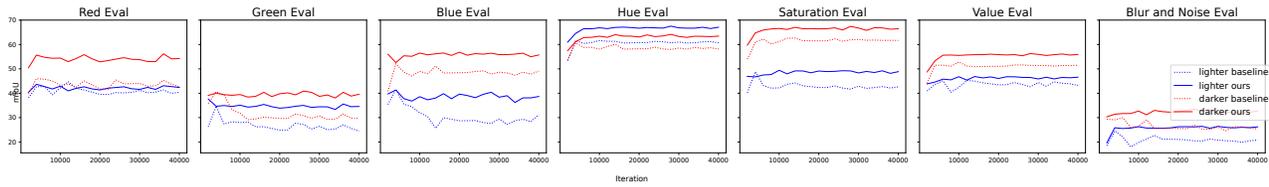


Figure 15. **Evaluation on perturbed test datasets over training iterations.** We show the evaluation on each perturbed dataset during training of our model and the baseline for VOC2012 dataset.

Q. Adaptive Sensitivity Analysis with Different Number of Levels

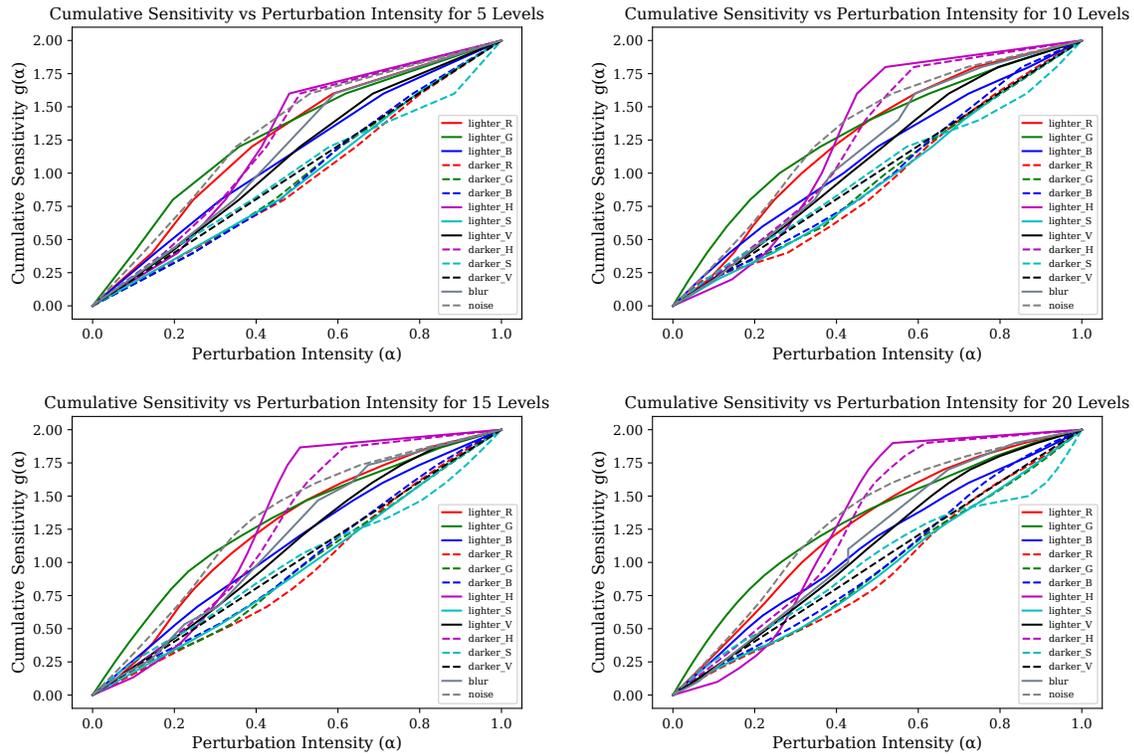


Figure 16. Visualization of cumulative sensitivity curve with varying number of levels L . We visualize the cumulative sensitivity curve in Equation 5 when computing for 5, 10, 15, and 20 levels. We find that even when we increase the number of levels, the curves remain *approximately* the same. Thus, we use 5 levels in our implementation to reduce compute for the sensitivity analysis step.