Gradient-Free Adversarial Training Against Image Corruption for Learning-based Steering

Anonymous Author(s)
Affiliation
Address
email

Abstract

We introduce a simple yet effective framework for improving the robustness of learning algorithm against (input) image corruptions for autonomous driving, due to both internal (e.g., sensor noises and hardware abnormalities) and external factors (e.g., lighting, weather, visibility, and other environmental effects). Using sensitivity analysis with FID-based parameterization, we propose a novel algorithm exploiting basis perturbations to improve the overall performance of autonomous steering and other image processing tasks, such as classification and detection, for self-driving cars. Our model not only improves the performance on the original dataset, but also achieves significant performance improvement on datasets with multiple and unseen perturbations, up to 87% and 77%, respectively. A comparative study drawn between our approach and other SOTA techniques confirms the effectiveness of our technique in improving the robustness of neural network training for learning-based steering.

1 Introduction

Autonomous driving is a complex task that requires many software and hardware components to operate reliably under highly disparate and often unpredictable conditions. In this work, we study steering as it contains both perception and control (two most critical components) for autonomous driving. While on the road, vehicles are likely to experience day and night, clear and foggy conditions, sunny and rainy days, as well as bright cityscapes and dark tunnels. All these external factors in conjunction with internal factors of the camera (e.g., those associated with hardware) can lead to quality variations in input data for image-based learning algorithms. One can harden machine learning systems to these degradations by simulating them at training time [6]. However, an algorithmic tool for analyzing the sensitivity of real-world neural network performance on the properties of training images is lacking. More importantly, a mechanism to leverage such a sensitivity analysis for improving learning outcomes needs to be developed. In this work, we quantify the influence of image quality on the task of "learning to steer," study how training on degraded and low-quality images can boost robustness to image corruptions, and provide a systematic approach to improve the performance of the learning algorithm based on quantitative analysis.

Image degradations can be simulated by varying attributes such as blur, noise, distortion, color representations (such as RGB or CMY) hues, saturation, and intensity values (HSV). However, identifying the correct combination of the simulated attribute parameters to obtain optimal performance on real data is a difficult—if not impossible—task, as it requires domain transfer and exploring a high dimensional parameterized space.

We then *design a systematic method for measuring the severity of image degradation and predicting the impact of such degradation on model performance*. Inspired by the use of image feature variance in sensitivity analysis [34], we measure the difference between real-world image distributions and

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

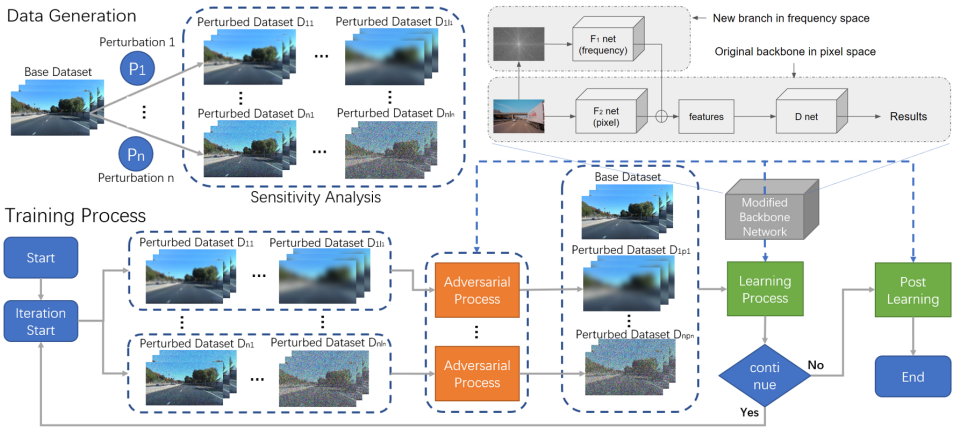

Figure 1: Pipeline of our method. Data generation: we generate perturbed datasets of each factor at multiple levels based on the FID-parameterized sensitivity analysis results. Training process: first stage: in each iteration, we first augment the training dataset with "adversarial images" given each perturbation; we then combine the base and perturbed datasets to train our model to maximize the overall performance. Iteration stops when break conditions are met. Second stage: in post-learning, the model is fine-tuned solely on clean data to boost accuracy, while performing special validation with an early break to maintain the performance on the perturbed data. A frequency-space branch is added to the backbone when frequency-related perturbations (e.g., blur, noise) need to be handled.

simulated/degraded image distributions using Fréchet Inception Distance (FID). Our results confirm that FID can help predict the performance of a model trained using simulated data and deployed in the real world. Next, we use FID between different simulated datasets as a unified metric to parameterize the severity of various image degradations due to different factors.

Borrowing concepts from the adversarial attack literature [28, 36, 44], we *build a scalable training scheme for enhancing the robustness of autonomous driving against various combinations of image degradations, while* increasing *the overall accuracy of the steering task on clean data*. Our proposed method constructs a dataset of adversarially degraded images by applying optimization within the space of possible degradations during training. As shown in Fig. 1, the method begins by training on a set of real and simulated/degraded images using arbitrary degradation parameters. During each training iteration, the parameters are updated to generate a new degradation set so that the model performance is (approximately) minimized. The network is then trained on these adversarially degraded images to promote robustness. A post-training is applied as the last step to further improve the performance on clean data without weakening robustness. Our proposed algorithm uses our FID-based parameterization to discretize the search space of degradation parameters and accelerates the process of finding optimal parameters.

Experiments show that our algorithm improves the performance of "learning to steer" up to **97%** in mean accuracy over baselines, and especially improves the performance on datasets contaminated with complex combinations of perturbations (up to **87%**). It additionally boosts the test performance on degradations that are not seen during training, including simulated snow, fog, and frost (up to **77%**). We also compare our approach with other SOTA techniques (e.g. data augmentation and adversarial training) on visual processing tasks such as detection and classification. Our method consistently achieves higher performance. In addition, our method is easy to implement and can be readily integrated with other frameworks like object detection, classification, regression, etc.

Furthermore, we propose a comprehensive robustness evaluation standard under four different scenarios: clean data, single-perturbation data, multi-perturbation data, and previously unseen data. While state-of-the-art studies usually conduct testing under one or two scenarios (e.g., ImageNet-C [20]), our work tests and verify results under four meaningful scenarios. We plan to release code and datasets for benchmarking "autonomous driving under perturbations" using, unseen factors like image corruptions in ImageNet-C [20], totaling **480** datasets and **26M** images.

2 Related Work

The influence of noise and distortion on real images for learning tasks has been well explored. For example, researchers have examined the impact of optical blur on convolutional neural networks and present a fine-tuning method for recovering lost accuracy using blurred images [42]. This

fine-tuning method resolves lost accuracy when images are distorted instead of blurred [51]. While these fine-tuning methods are promising, [11] find that tuning to one type of image quality reduction effect would cause poor generalization to other types of quality reduction effects. The comparison of image classification between deep neural networks and humans shows similar performance on good-quality images [12]. However, deep neural networks struggle significantly more than humans on low-quality, distorted, and noisy images. One study shows that adversarial perturbations are more prevalent in the Y channel in the YCbCr color space of images than the other two channels, while perturbations in RGB channels are equally distributed [29]. [43] studies the effect of Instagram filters on learning tasks, and introduces a lightweight de-stylization module that predicts parameters used for scaling and shifting feature maps to "undo" the changes incurred by filters.

Researchers have also explored how to improve the robustness of learning algorithms under various image quality degradations. One recent work [40] provides a novel Bayesian formulation for data augmentation. [9] proposes an approach to automatically search for improved data augmentation policies. [17] analyzes the performance of convolutional neural networks on quality degradations due to compression loss, noise, blur, and contrast, and introduces a method to improve the learning outcome. Another work [21] shows that self-supervision techniques can be used to improve model robustness and exceeds the performance of fully supervised methods. A recent method, AugMix (2019) [22] improves model robustness using data augmentation, where transformation compositions are used to create a new dataset that is visually and semantically similar to the original dataset. AugMix is compared to several other augmentation methods, which comprise Cutout [10], MixUp [46], and CutMix [45]. [14] proposes a technique to re-purpose software testing methods to augment the training data of DNNs, with the objective to improve model robustness. A recent work [18] improves model generalizability by first augmenting the training dataset with random perturbations, and then minimizing worst-case loss over the augmented data.

Our work differs from these studies in several regards. First, we simulate adversarial conditions of image factors instead of using commonplace image conditions. Second, we conduct a systematic sensitivity analysis for preparing datasets that are representative of image degradations from multiple factors at various levels. Third, our algorithm can work with the discretized parameter space while generalizing well to the continuous parameter space. Another advantage of our approach is that we can augment the training dataset without the derivatives of the factor parameters, which may not exist or are difficult to compute.

3 Background and Setup

Base Dataset. We choose four real-world driving corpuses as base datasets: Audi [16], Honda [31], Waymo [38], and SullyChen [7]. The Audi dataset is the most recent (2020); the Honda dataset has 100+ long-time driving videos; Waymo includes many environmental conditions such as weather and lightening and is a large dataset (390k frames for perception); and the SullyChen dataset focuses on the steering task and has the longest continuous driving image sequence without road branching. The details of these datasets are provided in Appendix A.3. Other datasets, such as CIFAR-100, for various computer vision tasks are also used to demonstrate the generalizability of our technique.

Basis Perturbation. We study nine basis perturbations: blur, noise, distortion, three-color (RGB) channels, and hues, saturation, and intensity values (HSV). Blur, noise, and distortion are among the most commonly used perturbations that can directly affect image quality. R, G, B, H, S, V channels are chosen because they are frequently used to represent image color spaces: RGB represent three basic color values of an image, while HSV represent three common metrics of an image. Other color spaces such as HSL or YUV have similar properties, hence are excluded.

We use Gaussian blur [2] (parameterized by standard deviation), additive white Gaussian noise (AWGN) (parameterized by standard deviation), and radial distortion [50] [49] (with radial distortion parameters $k1, k2$), respectively. For representing channel-level perturbations, we use a linear model: denote the value range of one channel C as $[a_C, b_C]$, in the darker direction, we set $v'_C = \alpha a_C + (1-\alpha)v_C$, while in the lighter direction, we set $v'_C = \alpha b_C + (1-\alpha)v_C$, where α is the severity parameter, v_C is the pixel value on clean image and v'_C is the perturbed pixel value. The default values are $a_C = 0$ and $b_C = 255$, with two exceptions. We set $a_V = 10$ to exclude a complete dark image on V channel, and $b_H = 179$ according to H channel definition. See examples in Fig. 2 and detailed description in Appendix. A.2.

Test Scenarios. While other studies usually test in only one or two scenarios [20], we test the performance of all methods in four scenarios with increasing complexity: Scenario 1: Base dataset.

Figure 2: Example images of quality degradation. Left: Five levels of quality reduction using the blur, noise, and distortion effects (three are shown). Right: The original images are shown in the middle row. The top row shows examples in the lighter direction per channel for R, G, B, H, S, V, while the bottom row shows examples in the darker direction per channel.

Test on the base dataset. Scenario 2: Single perturbation. Test on the datasets with perturbations, but each dataset only experiences one level of one perturbation (within blur, noise, distortion, R, G, ,B, H, S, V). Scenario 3: Combined perturbations. Test on the datasets with combinations of various perturbations at different levels, and each dataset has different levels of all perturbations. Scenario 4: Unseen perturbation. Test on the datasets with unseen perturbations at different levels. Specifically, we use "motion blur", "zoom blur", "pixelate", "jpeg compression", "snow", "frost", "fog" from ImageNet-C [20]. We apply the same perturbations/levels to all images within one dataset. (See details in Appendix A.1).

4 Improving Robustness of Learning-based Steering

In this section, we introduce our gradient-free adversarial training method to improve the robustness of learning-based steering using Sensitivity Analysis. Our method is among the first to *treat complex unforeseen perturbations as a functional combination of multiple simple "basis perturbations"*. As a result, the robustness of a model to several individual perturbations can result in robustness to combined or unseen perturbations. In addition, *via discretized sensitivity analysis*, our approach is the first to *use FID to enable cross-factor performance comparison*, i.e., making task sensitivity w.r.t. different perturbations comparable.

4.1 Basis Perturbations

There are various types of image corruptions due to different environmental effects or sensor variations, and often these corruptions can be approximated using a combination of basis perturbations, e.g., snow/frost in a scene may be approximated by a combination do water drop corruptions, Gaussian blur, and image whitening. Inspired by the concept of *basis function*, we explore whether it is possible to only training on a few "basis perturbations", resulting in robustness to complex perturbations. Our basis perturbation choice also contains "basis" representations of the color space in RGB and of image metrics in HSV which span color and intensity/brightness spaces, respectively. Blur and noise are designed to produce low and high frequency corruptions, while distortion captures the 2D positional transformation that may appear from small vibrations and motions of a camera. We also introduce two metrics to evaluate how well a training method can enable a model to perform on combined perturbations and previously unseen perturbations when learning only on single perturbations one at a time (see Sec. 5). Using basis perturbations to represent image corruption, we design an algorithm to train the network to better cope with degraded input image quality in autonomous driving.

4.2 Sensitivity Analysis

We use an adversarial training process in which we optimize corruption parameters to maximally disrupt model performance. If this were done using gradient optimization, it would require us to differentiate through the corruption operator. Not all gradients of basis perturbations can be easily derived, e.g., distortion. So, we introduce an efficient, simple, discretized gradient-free method. We use Sensitivity Analysis (SA) as the support for discretization and show the use of Fréchet Inception Distance (FID) [23] for cross-perturbation SA.

SA is commonly used to understand how the uncertainty in the model output (numeric or otherwise) can be apportioned to different sources of uncertainty in the model input [33] [34]. A recent work [48] shows that SA can be used to better interpret CNNs. Here, we use SA to study how distortions to model inputs (blur, noise, distortion, and RGB/HSV brightness shifts) can be apportioned to model output (i.e., MA of the predicted steering angle for the test data). We also use SA results to prepare datasets that are representative of image qualities at various degradation levels for model training.

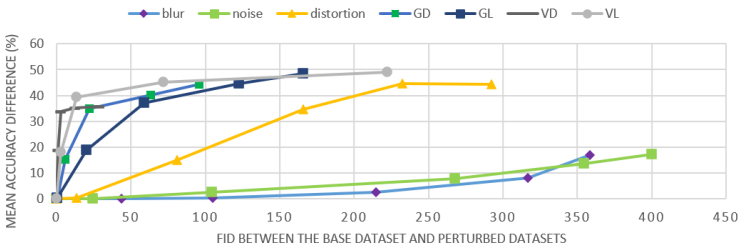

Figure 3: The relationship between FID and MA difference. GD/GL denotes G channel in darker/lighter direction, and VD/VL denotes V channel in darker/lighter direction, respectively. Sensitivity is represented by the first-order derivative of the curve.

We adopt the Fréchet Inception Distance (FID) [23] as a unified metric for our sensitivity analysis (instead of using the parameter values of each image factor). We explain three reasons for using FID and also a quantitative comparison with other metrics like L2 norm distance in Appendix A.4. In addition, we focus on the changes in model performance according to the changes in the input factors, i.e., FID. Thus we define the sensitivity as the first-order derivative of MA w.r.t FID:

$$sensitivity = \frac{\partial MA^*(R \bigoplus F(\mathbf{p}))}{\partial FID(R, R \bigoplus F(\mathbf{p}))},$$

where $MA^*(D)$ is the MA test result on dataset D with the model trained on the base dataset R, $FID(A, B)$ is the FID between datasets A and B, $F(\mathbf{p})$ is the perturbation with parameter $\mathbf{p}$ (e.g., the standard deviation of the Gaussian kernel), and $D \bigoplus F$ means the perturbed dataset by applying perturbation F to D.

Starting from empirically-selected parameters of each factor, we generate perturbed datasets and compute their corresponding MA using the trained model on R (see numeric results in Appendix A.8). We then map the correspondences between the MAs and the parameter values into the FID space. By leveraging this new MA-FID space, we can reduce and minimize the number of parameter samples for each perturbation, while the sampled dataset provides a similar curve to the original one to improve the computational efficiency of training (see Sec 4). Examples of the resulting images are shown in Fig. 2 (more in Appendix A.1). Detailed descriptions of the final perturbed datasets are provided in Appendix A.2. We also provide a detailed analysis of MA differences in FID space in Appendix A.4.

The final MA differences in the FID space are visualized in Fig. 3 (for blur, noise, distortion, and G and V channels; see the entire figure in Appendix A.4). Since FID aligns different factors into the same space, we can compare the performance loss of all factors at various levels. Notice that the values in the near-zero FID range (i.e., FID< 50) are more commonly found in real-world applications. We first observe that the learning-based steering task is more sensitive to the channel-level perturbations (i.e., R, G, B, H, S, V) than the image-level perturbations (i.e., blur, noise, distortion). Second, the task is least sensitive to blur and noise but most sensitive to the V channel, the intensity value. Third, for the same color channel, darker and lighter levels appear to have different MAs at the same FID values. Compared with other analysis studies on the "learning to steer" task, e.g., [39] and [47], our method is the first to transfer perturbations from multiple parameter spaces into one unified space to enable the cross-factor comparison, e.g., the task is more sensitive to the V-channel perturbation than perturbations in other image attributes.

4.3 Gradient-Free Adversarial Training

Our algorithm consists of two stages. In the first stage, we use an iterative min-max training process. At each iteration, we first choose one dataset from datasets (with different severities) of each basis perturbations (i.e., R, G, B, H, S, V, blur, noise, and distortion) that can minimize validation MA, then merge the 9 chosen datasets with the base dataset to train our model while maximizing MA. The first stage stops when a pre-specified number of iterations is reached or the MA loss is below a certain threshold. The second stage is designed to boost clean accuracy, where the model is fine-tuned solely on clean data. To meanwhile maintain the performance on the perturbed data, we terminate this stage if the overall validation accuracy on clean and perturbed data decreases. Otherwise, this stage continues until it reaches the maximum number of iteration or the MA loss decreases to our expected threshold. See more explanation and ablation study for the second stage in Appendix A.5. The design rationale of our architecture resembles adversarial training: we train the model to maximize accuracy

using the original dataset plus the perturbed datasets with the minimum accuracy in order to improve the robustness of the model. The loss function is the following:

$$\min_{\mathbf{p}} \max_{\theta} MA(\theta, U_{\mathbf{p}}),$$

where $\mathbf{p}$ represents a union of the parameter levels of all analyzed image factors, θ denotes the model parameters, $U_{\mathbf{p}}$ is the training dataset, and $\mathrm{MA}()$ is the function computing MA. Furthermore, we add a branch in the frequency space to the backbone network, since frequency space can help capture features of frequency-related perturbations such as blur and noise. The ablation study demonstrates the effectiveness of this branch's addition (see Table 1). Utilizing basis perturbation properties in Sec. 4.1, we perform adversarial training on single perturbations, but test on more complex cases, including combined perturbations and previously unseen perturbations. See Algorithm 1.

Algorithm 1: Improve robustness of learning-based steering

Result: a trained model parameterized by θ
Pre-processing:
Conduct sensitivity analysis and discretize the parameters of n factors into their corresponding $l_i, i = 1, \ldots, n$ levels
Generate new datasets for each factor with the discretized values from the base dataset R: $\{D_{i,j}\}, i = 1, 2, .., n, j = 1, 2, ..l_i$
Initialization:
Set $t = 0$, and initialize the maximum iterations T and the number of epochs k_1 and k_2
Initialize model parameters $\theta^{(0)}$
First stage:
while $t \leq T$ **do**
  For each factor, select the dataset D_{i,p_i} that can minimize the validation MA, where $p_i = \arg\min_j MA(\theta^{(t)}, D_{ij})$
  Merge all selected datasets $U_{\mathbf{p}} = (\bigcup_{i=1}^{n} D_{i,p_i}) \bigcup R$
  Train the network for k_1 epochs and update $\theta^{(t+1)}$ = train($\theta^{(t)}$, $U_{\mathbf{p}}$, k_1) to maximize $MA(\theta^{(t+1)}, U_{\mathbf{p}})$
  Break if stop conditions are met; otherwise set $t = t + 1$.
end
Second stage:
Train the network with $\theta^{(final)}$ = train($\theta^{(T)}$, R, k_2) and validate on $U_{\mathbf{p}}$ (for k_2 epochs or early break if conditions are met.)

Our method offers several advantages: 1) the training data is augmented without re-train the model, thus improving efficiency; 2) it provides the flexibility to generate datasets at various discretized levels of the factor parameters; 3) it does not require the derivatives of factor parameters; other methods that optimize factor parameters in the continuous space require computing derivatives that can be difficult (e.g. for distortion); 4) it can generalize to not only unseen parameters of individual factors but also the composition of unseen parameters of multiple factors; and 5) it is easy to implement and can be readily integrated with other frameworks like object detection, classification, regression, etc.

5 Experiments and Results

Setups: All experiments are conducted using one Intel(R) Xeon(TM) W-2123 CPU, two Nvidia GTX 1080 GPUs, and 32G RAM. We use the Adam optimizer [24] with learning rate 0.0001 and batch size 128 for training. The maximum number of epochs is 2,000. The datasets setup is shared with analysis experiments (see Sec. 3). We use the maximum MA improvement (MMAI), the average MA improvement (AMAI), and mean Corruption Errors (mCE) [20] as the evaluation metrics.

Backbone: We choose the model by [4] as the main backbone. We select this model as it has been used to steer an AV successfully in both real world [4] and virtual world [25]. In addition, six other networks are used to show the generalizability of our method.

Evaluation Metric: We use mean accuracy (MA) to evaluate the steering task since it represents overall performance under a variety of thresholds (similar to Mean Average Precision (mAP) for the detection task). We first define the accuracy with respect to a particular threshold τ as $acc_\tau = count(|v_{predicted} - v_{actual}| < \tau)/n$, where n denotes the number of test cases; $v_{predicted}$ and v_{actual}

indicate the predicted and ground-truth value, respectively. We compute MA as $\sum_\tau acc_{\tau \in \mathcal{T}}/|\mathcal{T}|$, where $\mathcal{T} = \{1.5, 3.0, 7.5, 15, 30, 75\}$ contains empirically selected thresholds of steering angles. Lastly, we use the maximum MA improvement (denoted as MMAI), the average MA improvement (denoted as AMAI), and mean Corruption Errors (mCE) [20] as the evaluation metrics.

	Scenarios									
	Clean	Single Perturbation			Combined Perturbation			Unseen Perturbation		
Method	AMAI↑	MMAI↑	AMAI↑	mCE↓	MMAI↑	AMAI↑	mCE↓	MMAI↑	AMAI↑	mCE↓
Data Augmentation	-0.44	46.88	19.97	51.34	**36.1**	11.97	75.84	27.5	7.92	81.51
Adversarial Training	-0.65	30.06	10.61	74.42	17.89	6.99	86.82	16.9	8.17	89.91
MaxUp	-7.79	38.30	12.83	66.56	26.94	16.01	72.60	23.43	5.54	81.75
AugMix	-5.23	40.27	15.01	67.49	26.81	15.45	68.38	28.70	8.85	87.79
Ours w/o FS	0.93	48.57	20.74	49.47	33.24	17.74	63.81	29.32	9.06	76.20
Ours	**2.12**	**49.97**	**23.92**	**37.30**	33.15	**22.12**	**54.38**	**33.16**	**13.81**	**61.61**

Table 1: Performance of different methods against the baseline [4] on SullyChen dataset. We compare the basic data augmentation method (simply combine all perturbed datasets into training), an adversarial training method [37], MaxUp [18], and AugMix [22]. Overall, our method outperforms all other methods (i.e., highest MA improvements and lowest error in mCEs) in practically all scenarios. Notice ours w/o FS (without frequency space) also outperforms other technique.

Comparison with different methods: we compare our method with four other methods: an adversarial training method [37], a basic data augmentation method, MaxUp [18], and AugMix [22], to see the performance improvement over the baseline method [4]. For the basic data augmentation method, we simply merge all perturbed datasets for model training.

From Table 1, we observe that our method outperforms other methods under all metrics in all scenarios: not only on the clean dataset but also on perturbed datasets. Notably, our algorithm improves the performance of "learning to steer" up to 50% in MMAI, while reducing mCE by 60% over the baseline (Scenario 2). Our method also improves the task performance using the *combined* datasets (Scenario 3) up to 33%. Lastly, when tested on unseen factors (Scenario 4), our algorithm maintains the best performance by 34% in MMAI, while reducing mCE to 62%.

Compared to AugMix [22], our adversarial approach can select the most challenging datasets for training, thus improving model robustness. MaxUp [18] selects only the worst case among all augmentation data, which may lead to the loss of data diversity. In contrast, our method selects the worst cases in *all* perturbation types (i.e., one dataset per factor), thus improving generalizability. Compared to [37], which performs an adversarial process on the entire pixel space with only norm constraints, our method is able to utilize vast prior information generated by sensitivity analysis to reduce the search space. Lastly, compared to the basic data augmentation method, which uses all generated data in training, our method selects the most useful data for training, and thus improves computational efficiency while minimizing data generation.

Comparison with different backbones: We also perform a comparison on three backbones: Nvidia network [4], Comma.ai network [35], and ResNet152 [19]. We conduct these experiments on the Honda dataset, without modification to the backbone network. The results shown in Table 2 indicate that our method outperforms AugMix in most cases. Generally, our method achieves better performance on shallow networks than deep networks. But even on very deep networks such as ResNet152, our method achieves more than 5% improvement in all cases, except Scenario 1.

Comparison on different datasets: To demonstrate that our method does not overfit a particular dataset, we experiment on three independent datasets: Audi [16], Honda [31], SullyChen [7], and Waymo [38]. We use the Nvidia network as the backbone for these experiments. Table 3 shows that our method can achieve consistently better performance across all three datasets. Furthermore, our method can obtain up to **97%, 87%, 77%** accuracy improvement on the single, combined, unseen perturbation, while achieving **90%, 74%, 61%** relative error reduction in some cases, respectively.

Performance on other image processing tasks: To show the generalizability of our method, we test our method on the classification task using CIFAR-100. For a fair comparison, we use the same setting as AugMix (i.e., same perturbations, training and testing data, backbone, and test code). As shown in Table 4, our method consistently outperforms all backbones significantly, reducing **16.3%** to **34.1%** in mean errors. The data for other methods come from the AugMix paper [22]. We also test our algorithm on the detection task in autonomous driving. We use the Audi dataset [16] and the Yolov4 network [3] as base settings, and implement ours based on Yolov4. The result shows that

	Scenarios									
	Clean	Single Perturbation			Combined Perturbation			Unseen Perturbation		
Method	AMAI↑	MMAI↑	AMAI↑	mCE↓	MMAI↑	AMAI↑	mCE↓	MMAI↑	AMAI↑	mCE↓
AugMix+Nvidia	-0.12	40.64	10.94	76.48	25.97	16.79	64.41	**22.23**	5.99	84.95
Ours+Nvidia	**2.48**	**43.51**	**13.51**	**67.78**	**28.13**	**17.98**	**61.12**	16.93	**6.70**	**80.92**
AugMix+Comma.ai	-5.25	55.59	9.56	86.31	31.32	**0.77**	**106.1**	37.91	7.97	89.99
Ours+Comma.ai	**0.36**	**62.07**	**15.68**	**70.84**	**38.01**	0.74	108.32	**42.54**	**12.15**	**77.08**
AugMix+ResNet152	-4.23	20.84	1.45	96.24	12.21	6.71	80.19	15.40	2.87	97.62
Ours+ResNet152	**-0.96**	**24.29**	**5.19**	**79.76**	**16.05**	**8.02**	**75.16**	**16.58**	**5.33**	**85.68**

Table 2: Performance improvement of different backbone networks against the baseline performance using the Honda dataset. Our method outperforms AugMix in most cases. Notice that the methods with ResNet152 do not improve as much as the first two networks because the ResNet152 baseline already has relatively high performance.

	Scenarios									
	Clean	Single Perturbation			Combined Perturbation			Unseen Perturbation		
Method	AMAI↑	MMAI↑	AMAI↑	mCE↓	MMAI↑	AMAI↑	mCE↓	MMAI↑	AMAI↑	mCE↓
AugMix on SullyChen	-5.23	40.27	15.01	67.49	26.81	15.45	68.38	28.70	8.85	87.79
Ours on SullyChen	**1.46**	**49.76**	**22.87**	**40.84**	**33.15**	**22.12**	**54.38**	**33.87**	**13.51**	**62.50**
AugMix on Audi	-8.24	81.89	32.22	55.27	75.49	50.23	41.98	73.06	27.39	77.51
Ours on Audi	**5.98**	**97.57**	**48.50**	**10.27**	**87.56**	**62.38**	**25.80**	**77.22**	**32.71**	**39.14**
AugMix on Honda10k	-0.12	40.64	10.94	76.48	25.97	16.79	64.41	**22.23**	5.99	84.95
Ours on Honda10k	**2.48**	**43.51**	**13.51**	**67.78**	**28.13**	**17.98**	**61.12**	16.93	**6.70**	**80.92**
AugMix on Honda100k	-11.41	63.85	14.08	70.64	**68.95**	47.69	40.12	**61.68**	16.32	88.36
Ours on Honda100k	**-2.55**	**67.35**	**19.88**	**53.26**	65.10	**48.60**	**36.94**	51.90	**18.29**	**72.84**
AugMix on Waymo	18.27	45.40	23.30	59.31	**22.95**	16.92	66.36	**57.65**	29.10	55.63
Ours on Waymo	**20.34**	**46.85**	**26.76**	**52.84**	21.34	**18.24**	**64.58**	56.98	**31.12**	**53.18**

Table 3: Performance improvement of different datasets with different sizes against the baseline using the Nvidia backbone. Our method outperforms AugMix considerably in most cases.

our algorithm improves robustness in most scenarios (about 3%-5% mAP on average). See detail in Sec. A.6.

Decoupling Ratio and Generalization Ratio: We propose two new metrics to evaluate how well a training method can allow a model to perform on combined perturbations and unseen perturbations respectively, while only learning on single perturbations one at a time.

Definition 5.1 ***(Decoupling Ratio)*** *Let $MA(D_1, D_2)$ be the mean accuracy (MA) function that outputs the MA result of a model trained on dataset D_1 and test on dataset D_2. Let $P_i^*(D)$ be the ith ($i = 1, ..., n$) perturbation taking dataset D as input and producing a perturbed dataset. Let $NP^*(D) = P_n^*(P_{n-1}^*(...P_1^*(D))...)$ (datasets with combined/nested perturbations), $UP^*(D) = \bigcup_{i=1}^{n} P_i^*(D)$ (union of datasets with single perturbation), D_{tr} be the training set, and D_{te} be the test set, then the decoupling ratio is*

$$r_d = \frac{MA(UP^*(D_{tr}), NP^*(D_{te}))}{MA(UP^*(D_{tr}), UP^*(D_{te}))}$$

$r_d = 1$ indicates that the model trained on data with the union of single perturbation images and test on combined perturbation images can perform as well as testing on a union of single perturbation images. Notice that different models and training methods may influence the MA function, thus may also influence the decoupling ratio.

Definition 5.2 ***(Generalization Ratio)*** *Let $MA(D_1, D_2)$ be the mean accuracy (MA) function that outputs the MA result of a model trained on dataset D_1 and test on dataset D_2. Let $P_i^*(D)$ be the ith ($i = 1, ..., n$) perturbation taking dataset D as input and producing a perturbed dataset, and $Q_j^*(D)$ be the jth ($j = 1, ..., m$) perturbation where $P_i^* \neq Q_j^*$. Let $UP^*(D) = \bigcup_{i=1}^{n} P_i^*(D)$ (union of datasets with single perturbation), $UQ^*(D) = \bigcup_{j=1}^{m} Q_j^*(D)$ (union of datasets with unseen single perturbation), D_{tr} be the training set, and D_{te} be the test set, then the decoupling ratio is*

$$r_g = \frac{MA(UP^*(D_{tr}), UQ^*(D_{te}))}{MA(UP^*(D_{tr}), UP^*(D_{te}))}$$

$r_g = 1$ indicates that the model trained on data containing the union of single perturbation image sets and tested on unseen perturbation images can perform as well as if it were tested on a union of

	Standard	Cutout	Mixup	CutMix	AutoAugment*	Adv Training	AUGMIX	Ours
AllConvNet	56.4	56.8	53.4	56.0	55.1	56.0	42.7	**25.6**
DenseNet	59.3	59.6	55.4	59.2	53.9	55.2	39.6	**26.3**
WideResNet	53.3	53.5	50.4	52.9	49.6	55.1	35.9	**20.5**
ResNeXt	53.4	54.6	51.4	54.1	51.3	54.4	34.9	**15.4**
Mean	55.6	56.1	52.6	55.5	52.5	55.2	38.3	**22.0**

Table 4: Average classification error (in %) across several architectures. Our method is able to reduce mean corruption errors than the previous SOTA methods by 16.3%-34.1% on CIFAR-100-C data.

known single-perturbation images. Similarly, different models and training methods may influence the MA function, thus may also influence the generalization ratio.

On our dataset ($UP^*(D_{tr})$ is the training set with our basis perturbations, $NP^*(D_{te})$ is the data in test scenario 3, and $UQ^*(D_{te})$ is the data in test scenario 4), our method can achieve a high decoupling ratio $r_d = \frac{0.7463}{0.8836} = 0.84$, and a high generalization ratio $r_d = \frac{0.8291}{0.8836} = 0.94$. We did not list these two ratios for other methods because they are not applicable to them (i.e. they are not trained on only single perturbations).

Effectiveness visualization: Using the salience map on several combined samples in Fig. 7 of Appendix A.7, we show that our method can help the network to focus on important areas (e.g., the road in front) instead of random areas on perturbed images.

We also show t-SNE [26] of feature embeddings from the baseline and proposed method in Fig. 8 of Appendix A.7. Features from our method are more uniformly distributed, indicating the reduction of the domain gaps created by the perturbations, thus improving robustness.

Benchmarking Datasets: We plan to release our driving datasets with perturbations for benchmarking. The benchmark will contain three collections, including Audi [16], Honda [31], Waymo [38], and SullyChen dataset [7]. Each of them will contain a base dataset, datasets with five levels of perturbation in blur, noise, and distortion, ten levels of variations in the channels R, G, B, H, S, V, multiple combined perturbations over all nine factors using our implementation, and five levels of each unseen simulated factor, including snow, fog, frost, motion blur, zoom blur, pixelate, and jpeg compression using ImageNet-C. There are 480 datasets and about 26M images in total. The ground-truth steering angles (or angular velocity of the car) for all images will also be provided for validation, along with the code of perturbed data generation (for improving the robustness of other learning tasks) and our code. We will blur the personally identifiable information. See the generation parameters in Appendix A.2.

6 Conclusion and Future Work

In this paper, we first analyze the influence of different image-quality attributes on the performance of the task "learning to steer". We study three image-level effects (i.e., blur, noise, and distortion) and six channel-level effects (i.e., R, G, B, H, S, V) and find that image degradations due to these effects can impact task performance at various degrees. By using FID as a unified metric, we conduct sensitivity analysis in the MA-FID space. Leveraging the sensitivity analysis results, we propose an effective and efficient training method to improve the generalization of learning-based steering under various image perturbations. Our model not only improves the task performance on the original dataset, but also achieves significant performance improvement on datasets with a mixture of perturbations (up to 87%), as well as unseen adversarial examples including snow, fog, and frost (up to 77%). These results show that our model is one of the most robust methods for the image-based autonomous driving system. Our method can be easily extended and applied beyond the set of factors and the learning algorithm analyzed in this paper. It can also generalize to analyzing any arbitrarily high number of image/input factors, other learning algorithms, and multimodal sensor data. Lastly, other autonomous systems where the *perception-to-control* functionality plays a key role can possibly benefit from our technique as well. We will release the datasets generated in this work for benchmarking the robustness of learning algorithms for autonomous driving, as well as the code.

Our method currently uses discretization to achieve efficient training, but further optimization for our implementation is possible. For example, the efficiency of our technique may be further improved using other methods, such as the reweighting strategy [32]. Also, the size of augmented data may be a limitation. Currently we generate data before training to speed up the training, but further exploration about how to minimize the size of augmented data is also a direction. In summary, we believe that our framework is generalizable to other image factors, learning algorithms, multimodal sensor data, and tasks in other application domains. These are all natural directions worth exploration.

References

[1] Siddharth Agarwal, Ankit Vora, Gaurav Pandey, Wayne Williams, Helen Kourous, and James McBride. Ford multi-av seasonal dataset. *arXiv preprint arXiv:2003.07969*, 2020.

[2] Isabelle Bégin and Frank Ferrie. Blind super-resolution using a learning-based approach. In *Proceedings of the 17th International Conference on Pattern Recognition (ICPR)*, volume 2, pages 85–89, 2004.

[3] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.

[4] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

[5] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.

[6] Qianwen Chao, Huikun Bi, Weizi Li, Tianlu Mao, Zhaoqi Wang, Ming C. Lin, and Zhigang Deng. A survey on visual traffic simulation: Models, evaluations, and applications in autonomous driving. *Computer Graphics Forum*, 39(1):287–308, 2019.

[7] Sully Chen. A collection of labeled car driving datasets, https://github.com/sullychen/driving-datasets, 2018.

[8] 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.

[9] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018.

[10] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.

[11] Samuel Dodge and Lina Karam. Quality resilient deep neural networks. *arXiv preprint arXiv:1703.08119*, 2017.

[12] Samuel Dodge and Lina Karam. A study and comparison of human and deep learning recognition performance under visual distortions. In *2017 26th international conference on computer communication and networks (ICCCN)*, pages 1–7. IEEE, 2017.

[13] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017.

[14] Xiang Gao, Ripon K Saha, Mukul R Prasad, and Abhik Roychoudhury. Fuzz testing based data augmentation to improve robustness of deep neural networks. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, pages 1147–1158, 2020.

[15] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.

[16] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S. Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, Tiffany Fernandez, Martin Jänicke, Sudesh Mirashi, Chiragkumar Savani, Martin Sturm, Oleksandr Vorobiov, Martin Oelker, Sebastian Garreis, and Peter Schuberth. A2D2: Audi Autonomous Driving Dataset. 2020.

[17] S. Ghosh, R. Shet, P. Amon, A. Hutter, and A. Kaup. Robustness of deep convolutional neural networks for image degradations. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2916–2920, 2018.

[18] Chengyue Gong, Tongzheng Ren, Mao Ye, and Qiang Liu. MaxUp: A Simple Way to Improve Generalization of Neural Network Training. *arXiv e-prints*, page arXiv:2002.09024, February 2020.

[19] 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.

[20] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.

[21] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[22] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *ICLR*, 2020.

[23] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *arXiv preprint arXiv:1706.08500*, 2017.

[24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[25] Weizi Li, David Wolinski, and Ming C. Lin. ADAPS: Autonomous driving via principled simulations. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 7625–7631, 2019.

[26] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

[27] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *The International Journal of Robotics Research (IJRR)*, 36(1):3–15, 2017.

[28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.

[29] Camilo Pestana, Naveed Akhtar, Wei Liu, David Glance, and Ajmal Mian. Adversarial Perturbations Prevail in the Y-Channel of the YCbCr Color Space. *arXiv e-prints*, page arXiv:2003.00883, February 2020.

[30] Matthew Pitropov, Danson Garcia, Jason Rebello, Michael Smart, Carlos Wang, Krzysztof Czarnecki, and Steven Waslander. Canadian adverse driving conditions dataset. *arXiv preprint arXiv:2001.10117*, 2020.

[31] Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, and Kate Saenko. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7699–7707, 2018.

[32] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. *arXiv preprint arXiv:1803.09050*, 2018.

[33] Andrea Saltelli. Sensitivity analysis for importance assessment. *Risk analysis*, 22(3):579–590, 2002.

[34] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.

[35] Eder Santana and George Hotz. Learning a driving simulator. *arXiv preprint arXiv:1608.01230*, 2016.

[36] Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *NeurIPS*, 2019.

[37] Manli Shu, Zuxuan Wu, Micah Goldblum, and Tom Goldstein. Preparing for the worst: Making networks less brittle with adversarial batch normalization. *arXiv preprint arXiv:2009.08965*, 2020.

[38] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. *arXiv*, pages arXiv–1912, 2019.

[39] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th international conference on software engineering*, pages 303–314, 2018.

[40] Toan Tran, Trung Pham, Gustavo Carneiro, Lyle Palmer, and Ian Reid. A bayesian data augmentation approach for learning deep models. In *Advances in neural information processing systems*, pages 2797–2806, 2017.

[41] 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.

[42] Igor Vasiljevic, Ayan Chakrabarti, and Gregory Shakhnarovich. Examining the impact of blur on recognition by convolutional networks. *arXiv preprint arXiv:1611.05760*, 2016.

[43] Zhe Wu, Zuxuan Wu, Bharat Singh, and Larry Davis. Recognizing instagram filtered images with feature de-stylization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:12418–12425, 04 2020.

[44] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L. Yuille, and Quoc V. Le. Adversarial examples improve image recognition. In *CVPR*, 2020.

[45] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *International Conference on Computer Vision (ICCV)*, 2019.

[46] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *International Conference on Learning Representations*, 2018.

[47] Mengshi Zhang, Yuqun Zhang, Lingming Zhang, Cong Liu, and Sarfraz Khurshid. Deeproad: Gan-based metamorphic testing and input validation framework for autonomous driving systems. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 132–142. IEEE, 2018.

[48] Ye Zhang and Byron Wallace. A sensitivity analysis of (and practitioners' guide to) convolutional neural networks for sentence classification. *arXiv preprint arXiv:1510.03820*, 2015.

[49] Zhengyou Zhang. A flexible new technique for camera calibration. *Technical Report MSR-TR-98-71, Microsoft Research*, 1998.

[50] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334, 2000.

[51] Yiren Zhou, Sibo Song, and Ngai-Man Cheung. On classification of distorted images with deep convolutional neural networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1213–1217. IEEE, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Sec. 6.
   (c) Did you discuss any potential negative societal impacts of your work? [No] Our method improves the robustness of target algorithms (methodology), and improve the safety for autonomous driving (application).
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sec.**??**.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Sec.**??**.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Data is provided as a dataset, see Sec. 5. Code will be provided after the paper is accepted.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Sec. 5, Appendix. A.2, and Appendix. A.3.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] We show experiment result on different backbones, different datasets, different tasks, compare to different methods, and perform consistently better. Also error bars for each of them will be too computationally expensive to get.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Sec. 5.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Sec. 3 and Sec. 5.
   (b) Did you mention the license of the assets? [No] We specify the source of the assets. People can find the asset's license in their official website.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We plan to provide a new benchmark, see Benchmarking Datasets subsection in Sec. 5.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] We use public data or code from their official website or git, no special data or code.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] See Benchmarking Datasets subsection in Sec. 5.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]