

Modeling Camera Effects to Improve Visual Learning from Synthetic Data

Alexandra Carlson*, Katherine A. Skinner*, Ram Vasudevan, and Matthew Johnson-Roberson

University of Michigan, Ann Arbor

Abstract. Recent work has focused on generating synthetic imagery to increase the size and variability of training data for learning visual tasks in urban scenes. This includes increasing the occurrence of occlusions or varying environmental and weather effects. However, few have addressed modeling variation in the sensor domain. Sensor effects can degrade real images, limiting generalizability of network performance on visual tasks trained on synthetic data and tested in real environments. This paper proposes an efficient, automatic, physically-based augmentation pipeline to vary sensor effects – chromatic aberration, blur, exposure, noise, and color temperature – for synthetic imagery. In particular, this paper illustrates that augmenting synthetic training datasets with the proposed pipeline reduces the domain gap between synthetic and real domains for the task of object detection in urban driving scenes.

Keywords: Deep learning, image augmentation, object detection.

1 Introduction

Deep learning has enabled impressive performance increases across a range of computer vision tasks. However, this performance improvement is largely dependent upon the size and variation of labeled training datasets that are available for a chosen task. For some tasks, benchmark datasets contain millions of hand-labeled images for the supervised training of deep neural networks (DNNs) [1, 2]. Ideally, we could compile a large, comprehensive training set that is representative of all domains and is labelled for all visual tasks. However, it is expensive and time-consuming to both collect and label large amounts of training data, especially for more complex tasks like detection or pixelwise segmentation [40]. Furthermore, it is practically impossible to gather a single real dataset that captures all of the variability that exists in the real world.

Two promising methods have been proposed to overcome the limitations of real data collection: graphics rendering engines and image augmentation pipelines. These approaches enable increased variability of scene features across an image set without requiring any additional manual data annotation. Recent work in rendering datasets has shown success in training DNNs with large amounts of

* Authors contributed equally to this work.




Fig. 1. Examples of object detection tested on KITTI for baseline unaugmented data (left) and for our proposed method (right). Blue boxes show correct detections; red boxes show detections missed by the baseline method but detected by our proposed approach for sensor-based image augmentation.

highly photorealistic, synthetic data and testing on real data [17, ?]. Pixel-wise labels for synthetic images can be generated automatically by rendering engines, greatly reducing the cost and effort it takes to create ground truth for different tasks. Recent work on image augmentation has focused on modeling environmental effects such as scene lighting, time of day, scene background, weather, and occlusions in training images as a way to increase the representation of these visual factors in training sets, thereby increasing robustness to these cases during test time [5, 6]. Another proposed augmentation approach is to increase the occurrence of objects of interest (such as cars or pedestrians) in images in order to provide more training examples of those objects in different scenes and spatial configurations [4, 7].

However, even with varying spatial geometry and environmental factors in an image scene, there remain challenges to achieving robustness of task performance when transferring trained networks between synthetic and real image domains. To further understand the gaps between synthetic and real datasets, it is worthwhile to consider the failure modes of DNNs in visual learning tasks. One factor that has been shown to contribute to degradation of performance and cross-dataset generalization for various benchmark datasets is sensor bias [8–11]. The interaction between the camera model and lighting in the environment can greatly influence the pixel-level artifacts, distortions, and dynamic range induced in each image [12–14]. Sensor effects, such as blur and overexposure, have been shown to decrease performance of object detection networks in urban driving scenes [15]. Examples of failure modes caused by over exposure, manifesting as missed detections, are shown in Figure 1. However, there still is an absence in the literature examining how to improve failure modes due to sensor effects for learned visual tasks in the wild.

In this work, we propose a novel framework for augmenting synthetic data with realistic sensor effects – effectively randomizing the sensor domain for syn-

thetic images. Our augmentation pipeline is based on sensor effects that occur in image formation and processing that can lead to loss of information and produce failure modes in learning frameworks – chromatic aberration, blur, exposure, noise and color cast. We show that our proposed method improves performance for object detection in urban driving scenes when trained on synthetic data and tested on real data, an example of which is shown in Figure 1. Our results demonstrate that sensor effects present in real images are important to consider for bridging the domain gap between real and simulated environments.

This paper is organized as follows: Section 2 presents related background work; section 3 details the proposed image augmentation pipeline; section 4 describes experiments and discusses results of these experiments and section 5 concludes the paper. Code for this paper can be found at <https://github.com/alexacarlson/SensorEffectAugmentation>.

2 Related Work

Domain randomization with synthetic data: Rendering and gaming engines have been used to synthesize large, labelled datasets that contain a wide variety of environmental factors that could not be feasibly captured during real data collection [3, 16]. Such factors include time of day, weather, and community architecture. Improvements to rendering engines have focused on matching the photorealism of the generated data to real images, which comes at a huge computational cost. Recent work on domain randomization seeks to bridge the reality gap by generating synthetic data with sufficient random variation over scene factors and rendering parameters such that the real data falls into this range of variation, even if rendered data does not appear photorealistic. Tobin et. al [27] focus on the task of object localization trained with synthetic data. They perform domain randomization over textures, occlusion levels, scene lighting, camera field of view, and uniform noise within the rendering engine, but their experiments are limited to highly simplistic toy scenes. Building on [27], Tremblay et al [?] generate a synthetic dataset via domain randomization for object detection of real urban driving scenes. They randomize camera viewpoint, light source, object properties, and introduce flying distractors. Our work focuses on image augmentation outside of the rendering pipeline and could be applied in addition to domain randomization in the renderer.

Augmentation with synthetic data: Shrivastava et al. recently developed SimGAN, a generative adversarial network (GAN) to augment synthetic data to appear more realistic. They evaluated their method on the the tasks of gaze estimation and hand pose estimation [19]. Similarly, Sixt et al. proposed RenderGAN, a generative network that uses structured augmentation functions to augment synthetic images of markers attached to honeybees [20]. The augmented images are used to train a detection network to track the honeybees. Both of these approaches focus on image sets that are homogeneously structured and low

resolution. We instead focus on the application of autonomous driving, which features highly varied, complex scenes and environmental conditions.

Traditional Augmentation Techniques: Standard geometric augmentations, such as rotation, translation, and mirroring, have become commonplace in deep learning for achieving invariance to spatial factors that are not relevant to the given task [24]. Photometric augmentations aim to increase robustness to differing illumination color and intensity in a scene. These augmentations induce small changes in pixel intensities that do not produce loss of information in the image. A well known example is the PCA-based color shift introduced by Krizhevsky et al. [1] to perform more realistic RGB color jittering. In contrast, our augmentations are modeled directly from real sensor effects and can induce large changes in the input data that mimics the loss of information that occurs in real data.

Sensor effects in learning: More generally, recent work has demonstrated that elements of the image formation and processing pipeline can have a large impact upon learned representation [28, 29, 10]. Andreopoulos and Tsotsos demonstrate the sensitivities of popular vision algorithms under variable illumination, shutter speed, and gain [8]. Doersch et al. show there is dataset bias introduced by chromatic aberration in visual context prediction and object recognition tasks [11]. They correct for chromatic aberration to eliminate this bias. Diamond et al. demonstrate that blur and noise degrade neural network performance on classification tasks [29]. They propose an end-to-end denoising and deblurring neural network framework that operates directly on raw image data. Rather than correcting for the effects of the camera during image formation of real images, we propose to augment synthetic images to simulate these effects. As many of these effects can lead to loss of information, correcting for them is non-trivial and may result in the hallucination of visual information in the restored image.

3 Sensor-based Image Augmentation

Figure 2 shows a side-by-side comparison of two real benchmark vehicle datasets, KITTI [38, 39] and Cityscapes [40], and two synthetic datasets, Virtual KITTI [16] and Grand Theft Auto [41, 17]. Both of the real datasets share many spatial and environmental visual features: both are captured during similar times of day, in similar weather conditions, and in cities regionally close together, with the camera located on a car pointing at the road. In spite of these similarities, images from these datasets are visibly different. This suggests that these two real datasets differ in their global pixel statistics. Qualitatively, KITTI images feature more pronounced effects due to blur and over-exposure. Cityscapes has a distinct color cast compared to KITTI. Synthetic datasets such as Virtual KITTI and GTA have many spatial similarities with real benchmark datasets, but are still visually distinct from real data. Our work aims to close the gap



Fig. 2. A comparison of images from the KITTI Benchmark dataset (upper left), Cityscapes dataset (upper right), Virtual KITTI (lower left) and Grand Theft Auto (lower right). Note that each dataset has differing color cast, brightness, and detail.

between real and synthetic data by modelling these sensor effects that can cause distinct visual differences between real world datasets. Figure 3 shows the archi-

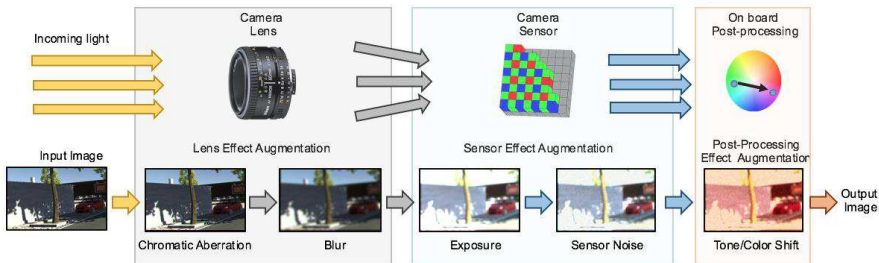


Fig. 3. A schematic of the image formation and processing pipeline used in this work. A given image undergoes augmentations that approximate the same pixel-level effects that a camera would cause in an image.

tecture of the proposed sensor-based image augmentation pipeline. We consider a general camera framework, which transforms radiant light captured from the environment into an image [30]. There are several stages that comprise the process of image formation and post-processing steps, as shown in the first row of Figure 3. The incoming light is first focused by the camera lens to be incident upon the camera sensor. Then the camera sensor transforms the incident light into RGB pixel intensity. On-board camera software manipulates the image (e.g., color space conversion and dynamic range compression) to produce the final output image. At each stage of the image formation pipeline, loss of information can occur to degrade the image. Lens effects can introduce visual distortions in an image, such as chromatic aberration and blur. Sensor effects can introduce over- or under-saturation depending on exposure, and high frequency pixel ar-

tifacts, based on characteristic sensor noise. Lastly, post-processing effects are implemented to shift the color cast to create a desirable output. Our image augmentation pipeline focuses on five total sensor effects augmentations to model loss of information that can occur at each stage during image formation and post-processing: chromatic aberration, blur, exposure, noise, and color shift. To model how these effects manifest in images in a camera, we implement the image processing pipeline as a composition of physically-based augmentation functions across these five effects, where lens effects are applied first, then sensor effects, and finally post-processing effects:

$$I_{aug.} = \phi_{color}(\phi_{noise}(\phi_{exposure}(\phi_{blur}(\phi_{chrom.ab.}(I)))))) \quad (1)$$

Note that these chosen augmentation functions are not exhaustive, and are meant to approximate the camera image formation pipeline. Each augmentation function is described in detail in the following subsections.

3.1 Chromatic Aberration

Chromatic aberration is a lens effect that causes color distortions, or fringes, along edges that separate dark and light regions within an image. There are two types of chromatic aberration, longitudinal and lateral, both of which can be modeled by geometrically warping the color channels with respect to one another [31]. Longitudinal chromatic aberration occurs when different wavelengths of light converge on different points along the optical axis, effectively magnifying the RGB channels relative to one another. We model this aberration type by scaling the green color channel of an image by a value S . Lateral chromatic aberration occurs when different wavelengths of light converge to the different points within the image plane. We model this by applying translations (t_x, t_y) to each of the color channels of an image. We combine these two effects into the following affine transformation, which is applied to each (x, y) pixel location in a given color channel C of the image:

$$\begin{bmatrix} x_C^{chrom.ab.} \\ y_C^{chrom.ab.} \\ 1 \end{bmatrix} = \begin{bmatrix} S & 0 & t_x \\ 0 & S & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_C \\ y_C \\ 1 \end{bmatrix} \quad (2)$$

3.2 Blur

While there are several types of blur that occur in image-based datasets, we focus on out-of-focus blur, which can be modeled using a Gaussian filter [33]:

$$G = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (3)$$

where x and y are spatial coordinates of the filter and σ is the standard deviation. The output image is given by:

$$I_{blur} = I * G \quad (4)$$

3.3 Exposure

To model exposure, we use the exposure density function developed in [34] [35]:

$$I = f(S) = \frac{255}{1 + e^{-A \times S}} \quad (5)$$

where I is image intensity, S indicates incoming light intensity, or exposure, and A is a constant value for contrast. We use this model to re-expose an image as follows:

$$S' = f^{-1}(I) + \Delta S \quad (6)$$

$$I_{exp} = f(S') \quad (7)$$

We vary ΔS to model changing exposure, where a positive ΔS relates to increasing the exposure, which can lead to over-saturation, and a negative value indicates decreasing exposure.

3.4 Noise

The sources of image noise caused by elements of the sensor array can be modeled as either signal-dependent or signal-independent noise. Therefore, we use the Poisson-Gaussian noise model proposed in [14]:

$$I_{noise}(x, y) = I(x, y) + \eta_{poiss}(I(x, y)) + \eta_{gauss} \quad (8)$$

where $I(x, y)$ is the ground truth image at pixel location (x, y) , η_{poiss} is the signal-dependent Poisson noise, and η_{gauss} is the signal-independent Gaussian noise. We sample the noise for each pixel based upon its location in a *GBRG* Bayer grid array assuming bilinear interpolation as the demosaicing function.

3.5 Post-processing

In standard camera pipelines, post-processing techniques, such as white balancing or gamma transformation, are nonlinear color corrections performed on the image to compensate for the presence of different environmental illuminants. These post-processing methods are generally proprietary and cannot be easily characterized [12]. We model these effects by performing translations in the CIELAB color space, also known as $L^*a^*b^*$ space, to remap the image tonality to a different range [36, 37]. Given that our chosen datasets are all taken outdoors during the day, we assume a D65 illuminant in our $L^*a^*b^*$ color space conversion.




Fig. 4. Example augmentations of GTA (left column) and VKITTI (right column) using the proposed sensor effect augmentation pipeline. Each image has a randomly sampled level of blur, chromatic aberration, exposure, sensor noise, and color temperature shift applied to it in an effort to model the visual structure/information loss caused by cameras when capturing real images.

3.6 Generating Augmented Training Data

The bounds on the sensor effect parameter regimes were chosen experimentally. The parameter selection process is discussed in more detail in Section 4. To augment an image, we first randomly sample from these visually realistic parameter ranges. Both the chosen parameters and the unaugmented image are then input to the augmentation pipeline, which outputs the image augmented with the camera effects determined by the chosen parameters. We augmented each image multiple times with different sets of randomly sampled parameters. Note that this augmentation method serves as a pre-processing step. Figure 4 shows sample images augmented with individual sensor effects as well as our full proposed sensor-based image augmentation pipeline. We use the original image labels as the labels for the augmented data. Pixel artifacts from cameras, like chromatic aberration and blur, make the object boundaries noisy. Thus, the original target labels are used to ensure that the network makes robust and accurate predictions in the presence of camera effects.

4 Experiments

We evaluate the proposed sensor-based image augmentation pipeline on the task of object detection on benchmark vehicle datasets to assess its effectiveness at bridging the synthetic to real domain gap. We apply our image augmentation pipeline to two benchmark synthetic vehicle datasets, each of which was rendered with different levels of photorealism. The first, Virtual KITTI (VKITTI) [16], features over 21000 images and is designed to model the spatial layout of KITTI

with varying environmental factors such as weather and time of day. The second is Grand Theft Auto (GTA) [41, 17], which features 21000 images and is noted for its high quality and increased photorealism compared to VKITTI. To evaluate the proposed augmentation method for 2D object detection, we used Faster R-CNN as our base network [42]. Faster R-CNN achieves relatively high performance on the KITTI benchmark test dataset, and many state-of-the-art object detection networks that improve upon these results use Faster R-CNN as their base architecture. For all experiments, we apply sensor effect augmentation pipeline to all images in the given dataset, then train an object detection network on the combination of original unaugmented data and sensor effect augmented data. We ran experiments to determine the number of sensor effect augmentations per image, and determined that optimal performance was achieved by augmenting each image in each dataset one time. To determine the bounds of the sensor effect parameter ranges from which to sample, we augmented small datasets of 2975 images by randomly sampling from increasingly larger parameter bounds and chose the ranges for each sensor effect that yielded the highest performance as well as visually realistic images. We found that the same parameter regime yielded optimal performance for both synthetic datasets. All of the trained networks are tested on a held out validation set of 1480 images from the KITTI training data and we report the Pascal VOC $AP50_{bbox}$ value for the car class. We also report the gain in $AP50_{bbox}$, which is the difference in performance relative to the baseline (unaugmented) dataset. We compare the performance of object detection networks trained on sensor-effect augmented data to object detection networks trained on unaugmented data as our baseline. For each dataset, we trained each Faster R-CNN network for 10 epochs using four Titan X Pascal GPUs in order to control for potential confounds between performance and training time.

4.1 Performance on baseline Object Detection Benchmarks

Table 1 shows results for FasterRCNN networks trained on unaugmented synthetic data and sensor-effect augmented data for both VKITTI and GTA. Note that we provide experiments trained on the full training datasets, as well as experiments trained on subsets of 2975 images to allow comparison of performance across differently sized datasets. Synthetic data augmented with the proposed method yields significant performance gains over the baseline (unaugmented) synthetic datasets. This is expected as, in general, rendering engines do not realistically model sensor effects such as noise, blur, and chromatic aberration as accurately as our proposed approach. Another important result for the synthetic datasets (both VKITTI and GTA), is that, by leveraging our approach, we are able to outperform the networks trained on over 20000 unaugmented images with a tiny subset of 2975 images augmented with using our approach. This means that not only can networks be trained faster but also when training with synthetic data, varying camera effects can outweigh the value of simply generating more data with varied spatial features. The VKITTI baseline dataset tested on KITTI performs relatively well compared to GTA, even though GTA

Table 1. Object detection trained on synthetic data, tested on KITTI

Virtual KITTI		
Training Set	AP_{car}	Gain
2975 Baseline	54.60	—
2975 Prop. Method	61.88	↑ 7.28
Full Baseline (21K)	58.25	—
Full Prop. Method	62.52	↑ 4.27
—		
GTA		
Training Set	AP_{car}	Gain
2975 Baseline	46.83	—
2975 Prop. Method	51.24	↑ 4.41
Full Baseline (21K)	49.80	—
Full Baseline (50K)	53.26	—
Full Prop. Method	55.85	↑ 6.05

is a more photorealistic dataset. This can most likely be attributed to the similarity in spatial layout and image features between VKITTI and KITTI. With our proposed approach, VKITTI gives comparable performance to the network trained on the Cityscapes baseline, showing that synthetic data augmented with our proposed sensor-based image pipeline can perform comparably to real data for cross-dataset generalization.

4.2 Comparison to other Augmentation Techniques

We ran experiments to compare our proposed method to photometric augmentation, specifically PCA-based color shift [1], complex spatial/geometric augmentations, specifically elastic deformation [47], standard additive gaussian noise augmentation, and a suite of standard spatial augmentations, specifically random rotations, scaling, translations, and cropping. We provide the results of training Faster-RCNN networks on the full VKITTI and GTA datasets augmented with the above methods in Table 2. All networks were tested on the same held-out set of KITTI images as used in the previous object detection experiments. Our results show that our proposed method drastically outperforms the other standard augmentation techniques, and that for certain synthetic data, spatial augmentations actually decrease performance on real data. This suggests that the proposed sensor effect augmentations capture more salient visual structure than traditional, non-photorealistic augmentation methods. We hypothesize this is because the physically-based sensor augmentations better model the information loss and the resulting global pixel-statistics that occur in real images. For example, our proposed method uses LAB space color transformation to alter the

color cast of an image, where as traditional approaches use RGB space. LAB space is device independent, so it results in a more accurate, physically-based augmentation than [1].

Table 2. We provide the results of training Faster-RCNN networks on GTA and Virtual KITTI augmented with various augmentation methods. All networks were tested on KITTI.

Virtual KITTI		
Augmentation Method	AP_{Car}	Gain
Baseline	58.25	—
Prop. Method	62.52	↑ 4.27
<i>Krishevsky et. al [1]</i>	59.09	↑ 0.84
<i>Ronneberger et. al [47]</i>	56.56	↓ 1.69
Additive Gaussian Noise	56.98	↓ 1.27
Random Rotation, Scale, Transl., Crop	55.11	↓ 3.14
GTA		
Augmentation Method	AP_{Car}	Gain
Baseline (21k)	49.80	—
Prop. Method (21k)	55.85	↑ 6.05
<i>Krishevsky et. al [1]</i>	51.62	↑ 1.88
<i>Ronneberger et. al [47]</i>	48.94	↑ 0.14
Additive Gaussian Noise	52.01	↑ 2.21
Random Rotation, Scale, Transl., Crop	50.11	↑ 0.31

4.3 Ablation Study

To evaluate the contribution of each sensor effect augmentation on performance, we used the proposed pipeline to generate datasets with only one type of sensor effect augmentation. We trained Faster-RCNN on each of these datasets augmented with single augmentation functions, the results of which are given in Table 3. Performance increases across all ablation experiments for training on synthetic data. This further validates our hypothesis that each of the sensor effects are important for closing the gap between synthetic and real data.

4.4 Failure Mode Analysis

Figure 5 shows the qualitative results of failure modes of FasterRCNN trained on each synthetic training dataset and tested on KITTI, where the blue bounding box indicates correct detections and the red bounding box indicate a missed

Table 3. Ablation study for object detection trained on synthetic data, tested on KITTI

Virtual KITTI			
Training Set	Augmentation Type	AP_{car}	Gain
2975 Baseline	<i>None</i>	54.60	—
2975 Prop. Method	Chrom. Ab.	61.08	↑ 6.48
2975 Prop. Method	Blur	59.72	↑ 5.12
2975 Prop. Method	Exposure	57.37	↑ 2.77
2975 Prop. Method	Sensor Noise	58.60	↑ 4.00
2975 Prop. Method	Color Shift	58.59	↑ 3.99
GTA			
Training Set	Augmentation Type	AP_{car}	Gain
2975 Baseline	<i>None</i>	46.83	—
2975 Prop. Method	Chrom. Ab.	48.92	↑ 2.09
2975 Prop. Method	Blur	49.17	↑ 2.34
2975 Prop. Method	Exposure	47.95	↑ 1.12
2975 Prop. Method	Sensor Noise	48.09	↑ 1.26
2975 Prop. Method	Color Shift	48.61	↑ 1.78

detection for the baseline that was correctly detected by our proposed method. Qualitatively, it appears that our method more reliably detects instances of cars that are small in the image, in particular in the far background, at a scale in which the pixel statistics of the image are more pronounced. Note that our method also improves performance on car detections for cases where the image is over-saturated due to increased exposure, which we are directly modeling through our proposed augmentation pipeline. Additionally, our method produces improved detections for other effects that obscure the presence of a car, such as occlusion and shadows, even though we do not directly model these effects. This may be attributed to increased robustness to effects that lead to loss of visual information about an object in general.

5 Conclusions

We have proposed a novel sensor-based image augmentation pipeline for augmenting synthetic training data input to DNNs for the task of object detection in real urban driving scenes. Our augmentation pipeline models a range of physically-realistic sensor effects that occur throughout the image formation and post-processing pipeline. These effects were chosen as they lead to loss of information or distortion of a scene, which degrades network performance on learned visual tasks. By training on our augmented datasets, we can effectively increase dataset size and variation in the sensor domain, without the need for



Fig. 5. Virtual KITTI examples are in the left column, GTA examples are in the right column. Blue boxes show correct detections; red boxes show detections missed by the FasterRCNN network trained on baseline, unaugmented image datasets but detected by FasterRCNNs trained on data augmented using our proposed approach for sensor-based image augmentation.

further labeling, in order to improve robustness and generalizability of resulting object detection networks. We achieve significantly improved performance across a range of benchmark synthetic vehicle datasets, independent of the level of photorealism. Overall, our results reveal insight into the importance of modeling sensor effects for the specific problem of training on synthetic data and testing on real data.

Acknowledgements

This work was supported by a grant from Ford Motor Company via the Ford-UM Alliance under award N022884, and by the National Science Foundation under Grant No. 1452793.

References

1. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems. (2012) 1097–1105

2. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
3. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2016) 3234–3243
4. Alhaija, H.A., Mustikovela, S.K., Mescheder, L., Geiger, A., Rother, C.: Augmented reality meets deep learning for car instance segmentation in urban scenes. In: *Proceedings of the British Machine Vision Conference*. Volume 3. (2017)
5. Zhang, H., Sindagi, V., Patel, V.M.: Image de-raining using a conditional generative adversarial network. *arXiv preprint arXiv:1701.05957* (2017)
6. Veeravasaru, V., Rothkopf, C., Visvanathan, R.: Adversarially tuned scene generation. *arXiv preprint arXiv:1701.00405* (2017)
7. Huang, S., Ramanan, D., undefined, undefined, undefined, undefined: Expecting the unexpected: Training detectors for unusual pedestrians with adversarial imposters. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* **00** (2017) 4664–4673
8. Andreopoulos, A., Tsotsos, J.K.: On sensor bias in experimental methods for comparing interest-point, saliency, and recognition algorithms. Volume 34., *IEEE* (2012) 110–126
9. Song, S., Lichtenberg, S.P., Xiao, J.: Sun rgb-d: A rgb-d scene understanding benchmark suite. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2015) 567–576
10. Dodge, S., Karam, L.: Understanding how image quality affects deep neural networks. In: *Quality of Multimedia Experience (QoMEX), 2016 Eighth International Conference on, IEEE* (2016) 1–6
11. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2015) 1422–1430
12. Grossberg, M.D., Nayar, S.K.: Modeling the space of camera response functions. Volume 26., *IEEE* (2004) 1272–1282
13. Couzinie-Devy, F., Sun, J., Alahari, K., Ponce, J.: Learning to estimate and remove non-uniform image blur. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2013) 1075–1082
14. Foi, A., Trimeche, M., Katkovnik, V., Egiazarian, K.: Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. Volume 17., *IEEE* (2008) 1737–1754
15. Ramanagopal, M.S., Anderson, C., Vasudevan, R., Johnson-Roberson, M.: Failing to learn: Autonomously identifying perception failures for self-driving cars. *CoRR* **abs/1707.00051** (2017)
16. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2016) 4340–4349
17. Johnson-Roberson, M., Barto, C., Mehta, R., Sridhar, S.N., Rosaen, K., Vasudevan, R.: Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? In: *Robotics and Automation (ICRA), 2017 IEEE International Conference on, IEEE* (2017) 746–753
18. Sankaranarayanan, S., Balaji, Y., Jain, A., Lim, S., Chellappa, R.: Unsupervised domain adaptation for semantic segmentation with gans. *CoRR* **abs/1711.06969** (2017)

19. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. arXiv preprint arXiv:1612.07828 (2016)
20. Sixt, L., Wild, B., Landgraf, T.: Rendergan: Generating realistic labeled data. arXiv preprint arXiv:1611.01331 (2016)
21. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint arXiv:1703.10593 (2017)
22. Eitel, A., Springenberg, J.T., Spinello, L., Riedmiller, M., Burgard, W.: Multimodal deep learning for robust rgb-d object recognition. In: Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, IEEE (2015) 681–687
23. Wu, R., Yan, S., Shan, Y., Dang, Q., Sun, G.: Deep image: Scaling up image recognition. arXiv preprint arXiv:1501.02876 7(8) (2015)
24. Hauberg, S., Freifeld, O., Larsen, A.B.L., Fisher, J., Hansen, L.: Dreaming more data: Class-dependent distributions over diffeomorphisms for learned data augmentation. In: Artificial Intelligence and Statistics. (2016) 342–350
25. Paulin, M., Revaud, J., Harchaoui, Z., Perronnin, F., Schmid, C.: Transformation pursuit for image classification. In: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE (2014) 3646–3653
26. Kim, H.E., Lee, Y., Kim, H., Cui, X.: Domain-specific data augmentation for on-road object detection based on a deep neural network. In: Intelligent Vehicles Symposium (IV), 2017 IEEE, IEEE (2017) 103–108
27. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on, IEEE (2017) 23–30
28. Kanan, C., Cottrell, G.W.: Color-to-grayscale: does the method matter in image recognition? *PLoS one* 7(1) (2012) e29740
29. Diamond, S., Sitzmann, V., Boyd, S., Wetzstein, G., Heide, F.: Dirty pixels: Optimizing image classification architectures for raw sensor data. (2017)
30. Karaimer, H.C., Brown, M.S.: A software platform for manipulating the camera imaging pipeline. In: European Conference on Computer Vision, Springer (2016) 429–444
31. Kang, S.B.: Automatic removal of chromatic aberration from a single image. In: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE (2007) 1–8
32. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in Neural Information Processing Systems. (2015) 2017–2025
33. Cheong, H., Chae, E., Lee, E., Jo, G., Paik, J.: Fast image restoration for spatially varying defocus blur of imaging sensor. *Sensors* 15(1) (2015) 880–898
34. Bhukhanwala, S.A., Ramabadran, T.V.: Automated global enhancement of digitized photographs. *IEEE Transactions on Consumer Electronics* 40(1) (Feb 1994) 1–10
35. Messina, G., Castorina, A., Battiato, S., Bosco, A.: Image quality improvement by adaptive exposure correction techniques. In: Multimedia and Expo, 2003. ICME '03. Proceedings. 2003 International Conference on. Volume 1. (July 2003) 1–549–52 vol.1
36. Hunter, R.S.: Accuracy, precision, and stability of new photoelectric color-difference meter. In: *Journal of the Optical Society of America*. Volume 38. (1948) 1094–1094

37. Annadurai, S.: Fundamentals of digital image processing, Pearson Education India (2007)
38. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE (2012) 3354–3361
39. Fritsch, J., Kuehnl, T., Geiger, A.: A new performance measure and evaluation benchmark for road detection algorithms. In: International Conference on Intelligent Transportation Systems (ITSC). (2013)
40. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 3213–3223
41. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In Leibe, B., Matas, J., Sebe, N., Welling, M., eds.: European Conference on Computer Vision (ECCV). Volume 9906 of LNCS., Springer International Publishing (2016) 102–118
42. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. (2015) 91–99
43. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (June 2015)
44. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR **abs/1409.1556** (2014)
45. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: The IEEE International Conference on Computer Vision (ICCV). Volume 2. (2017) 6
46. Hoffman, J., Wang, D., Yu, F., Darrell, T.: Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. CoRR **abs/1612.02649** (2016)
47. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention, Springer (2015) 234–241