

ConstScene: A Dataset and Model for Advancing Robust Semantic Segmentation in Construction Environment

Maghsood Salimi¹, Mohammad Loni², Sara Afshar², Antonio Cicchetti¹, and Marjan Sirjani¹

¹ School of Innovation, Design and Engineering, Mälardalen University, Sweden

² Future Solutions Department, Volvo Construction Equipment, Sweden
{maghsood.salimi, antonio.cicchetti, marjan.sirjani}@mdu.se,
{mohammad.loni, sara.afshar}@volvo.com

Abstract. The increasing demand for autonomous machines in construction environments necessitates the development of robust object detection algorithms that can perform effectively across various weather and environmental conditions. This paper introduces a new semantic segmentation dataset specifically tailored for construction sites, taking into account the diverse challenges posed by adverse weather and environmental conditions. The dataset is designed to enhance the training and evaluation of object detection models, fostering their adaptability and reliability in real-world construction applications. Our dataset comprises annotated images captured under a wide range of different weather conditions, including but not limited to sunny days, rainy periods, foggy atmospheres, and low-light situations. Additionally, environmental factors such as the existence of dirt/mud on the camera lens are integrated into the dataset through actual captures and synthetic generation to simulate the complex conditions prevalent in construction sites. We also generate synthetic images of the annotations including precise semantic segmentation masks for various objects commonly found in construction environments, such as wheel loader machines, personnel, cars, and structural elements. To demonstrate the dataset’s utility, we evaluate state-of-the-art object detection algorithms on our proposed benchmark. The results highlight the dataset’s success in adversarial training models across diverse conditions, showcasing its efficacy compared to existing datasets that lack such environmental variability.

Keywords: Dataset · Robust Object Detection · Adversarial Attacks · Semantic Segmentation · Construction Environment

1 Introduction

The growing need for autonomous machines in construction environments underscores the importance of advancing resilient object detection algorithms [13, 19, 24]. To ensure the applicability of the object detection model to unforeseen

situations, the training dataset should include diverse and representative data encompassing various weather and lighting conditions [1,17,25]. Meeting this demand is crucial for enhancing the reliability of autonomous driving in the heavy machinery industry.

However, the scarcity of sufficient training data including diverse environmental conditions from construction environments hinders the ability of object detection models to generalize effectively. In addition, it can exacerbate issues like overfitting, where models become overly specialized to the limited data available. Addressing the problem of inadequate training data is pivotal for developing robust and versatile object detection models capable of meeting the demands of complex and dynamic environments.

To address these challenges, this paper introduces a large-scale semantic segmentation dataset designed exclusively for construction sites, considering the varied challenges presented by diverse weather conditions and situations occurring during construction operations. We call this dataset robust semantic segmentation in construction environments, or in short: **ConstScene**. This dataset is designed to enhance the training and evaluation of semantic segmentation models, fostering their adaptability and reliability in real-world construction applications.

Contributions. **ConstScene** comprises annotated images captured under a wide range of weather conditions, including but not limited to sunny days, rainy periods, foggy atmospheres, and low-light situations. Furthermore, environmental factors such as dust, and the existence of dirt or mud on the camera lens are systematically integrated into the dataset to simulate the complex conditions prevalent in construction environments. The annotations include precise semantic segmentation masks for objects commonly found in construction environments, such as wheel loaders, personnel, cars, and structural elements. We augment our dataset by (i) applying the blurring effect; (ii) adding noise; and (iii) removing random regions through a cut-out process simulating the presence of dirt or mud on the camera lens. Finally, employing a monocular depth estimation model, we provide a depth map for every image in our dataset (Section 3.4).

To demonstrate the utility of our dataset, we evaluate two state-of-the-practice and state-of-the-art object detection algorithms, U-Net [21] and SegFormer [29], by training on the **ConstScene** dataset. According to our experiments, the test accuracy of U-Net and SegFormer trained on the non-augmented **ConstScene** dataset is up to 55.48% and 82.27%, respectively. The dataset and the code for training models and data augmentation are made available on the GitHub repository through: <https://github.com/RobustInsight/ConstScene/>.

2 Related Work

To the best of our knowledge, **ConstScene** is the first high-resolution public semantic segmentation dataset captured from a construction environment considering a wide range of objects and environmental conditions. This study also

sheds light on the impact of data augmentation on the segmentation model’s effectiveness in handling perturbed scenes.

On-Road Vision Datasets. Several research works have introduced vision datasets designed for applications in on-road autonomous driving [3, 4, 6, 10, 12, 31]. While earlier studies have achieved success in collecting extensive datasets with significant variations, they overlooked the inclusion of construction sites. In other words, they assumed that autonomous machines would not navigate through construction areas.

Off-Road Vision Datasets. [14] focused on guardrail detection in construction sites using a VGG16 model trained on a synthetic dataset containing 6,000 images. [26] collected images of five different construction machines, including loaders, bulldozers, excavators, backhoe diggers, and rollers. The MOCS dataset [30] gathered more than 40,000 images from various construction sites for the annotation of 13 distinct types of moving objects. Employing pixel segmentation, they annotated the objects and evaluated their performance on various Convolutional Neural Networks (CNN). ACID [28] collected 10,000 construction machine images for testing various object detection algorithms. SODA dataset [8] is a comprehensive object detection dataset for construction sites. SODA was released with more than 20k images encompassing four object categories including worker, material, machine, and layout.

Table 1 compares **ConstScene** with off-road autonomous driving datasets. Compared to prior studies, **ConstScene** is the first to address semantic segmentation in challenging environments, such as rainy, dusty, and foggy. Plus, we employ data augmentation techniques to robustify the model training against perturbed data. Notably, our dataset is openly accessible, ensuring result reproducibility. This facilitates researchers to focus on introducing novel detection techniques rather than repeating the time-consuming data collection process.

Table 1: Summary of off-road autonomous driving datasets. All methods utilized a monocular RGB camera for data collection.

Dataset	Year	Task	Static & Moving Objects	Challenging Weather	Data Perturbation	Availability
Kolar et al. [14]	2018	Object Detection	✗	✗	✗	✗
MOCS [30]	2021	Semantic Segmentation	✗	✓	✗	✗
ACID [28]	2021	Object Detection	✗	✓	✗	✗
SODA [8]	2022	Object Detection	✓	✓	✗	✗
ConstScene (Ours)	2024	Semantic Segmentation	✓	✓	✓	✓

3 Data Collection and Preparation

3.1 Data Collection and Sensor Setup

The selection of objects in image datasets is vital for training models to precisely recognize elements such as human, machines, and piles, ultimately enhancing

the reliability and decision-making capabilities of self-driving machines in diverse and dynamic construction environments. In this paper, we try to include commonly used objects in construction environments, including construction machine (wheel-loader), crusher, human, pile, road, and background regions.

The image data collection process poses several challenges, to include variability in the dataset such as lighting and weather conditions, diverse backgrounds, occlusions, and the need to include authentic situations such as existing dust on the camera lens or workers climbing the ladder of a heavy machine. To address these challenges, we carefully have selected representative samples considering diverse environment conditions to ensure model generalization. Plus, we employ data augmentation (Section 3.3), to artificially manipulate the input images to simulate variations in weather conditions. Occlusion challenges are mitigated by collecting images from multiple angles, while meticulous labeling and annotation processes are undertaken to provide the necessary ground truth for training robust models. For collecting image data, we utilized an RGB camera featuring a 1280×720 resolution, capturing images in JPEG format at a rate of 15 Frame-Per-Second.

3.2 Data Annotation

We leverage Roboflow [9] tool to annotate images. The tool allows us to label objects with semantic segmentation masks, enabling precise annotations for object classes. For each image in the dataset, a single Portable Network Graphics (PNG) mask is generated to show the segmented regions corresponding to different classes, providing a representation of the pixel-level annotation. After labeling the whole dataset, images were randomly divided into three categories train, validation, and test.

3.3 Data Augmentation

Image augmentation plays a crucial role in enhancing the robustness and generalization capabilities of semantic segmentation models [2]. By diversifying the training dataset through changes in contrast, and/or adding perturbations using noise, ultimately we improve the model’s ability to perform well on unseen or slightly different data during testing. In this paper, we leverage three data augmentation techniques by (i) blurring images through taking 25 neighbor pixels and averaging them; (ii) adding pepper-and-salt noise to 1% of pixels; and (iii) removing eight random regions with the size of 3% of image size through a cut-out process for imitating dirty and muddy camera lens. In the first two methods the added perturbation may not be visible to human eye but still can make the model to miss-classify. These augmentation techniques have been applied to the images of the train set in the dataset. Fig. 1 shows several samples from the dataset where these three perturbation techniques are applied.

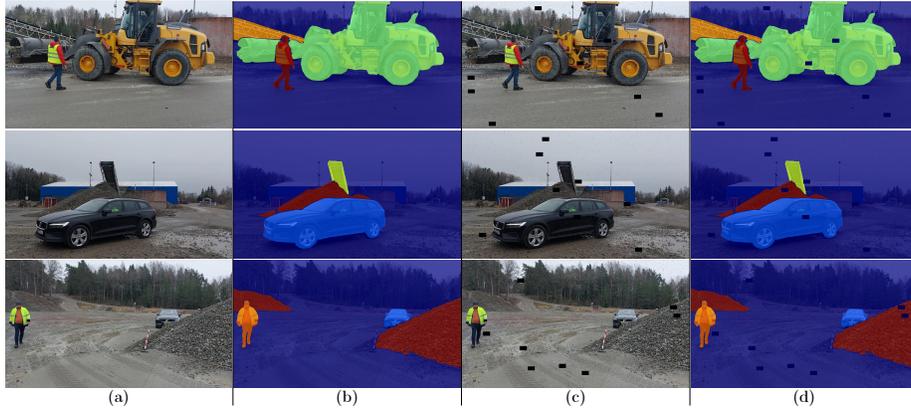


Fig. 1: (a) Original image. (b) Overlay mask of the original image. (c) Augmented image. (d) Overlay mask of the augmented image.

3.4 Depth Estimation

Depth estimation is the task of determining the distance from the camera to objects in a scene [15,16,18]. Depth estimation is crucial for various applications, such as robotics, autonomous machines, and augmented/virtual reality where environment perception is required for decision making. In autonomous machines application, depth estimation plays a crucial role for obstacle avoidance that contributes to safety. Given the significance of depth estimation in autonomous driving, we provide depth maps of images in the dataset using a monocular depth estimation method [20]. Fig. 2 shows multiple RGB samples of `ConstScene` with corresponding depth maps across diverse weather and working conditions. The red circles in Fig. 2.c to Fig. 2.e mark the areas of disparity with low-confidence matches for conditions with the presence of fog, dust, and dirt on the camera lens. As can be seen, the impact of adverse weather conditions significantly diminishes the accuracy of depth estimation, emphasizing the crucial need to address this challenge in computer vision pipelines employed in construction environments.

4 Dataset Statistics

`ConstScene` is available in two distinct versions: the original and the augmented version. The original version comprises 3470 images as training, validation, and test sets, while the augmented version consists of 6240 images. The partition ratio of the training, validation, and test sets for the original dataset are 80%, 10%, and 10%, respectively. On the other hand, for the augmented version, these sets were allocated at proportions of 90%, 5%, and 5%. Fig. 3 illustrates the class distribution in the dataset, indicating the number of images per class across the training, validation, and test sets in the original dataset.

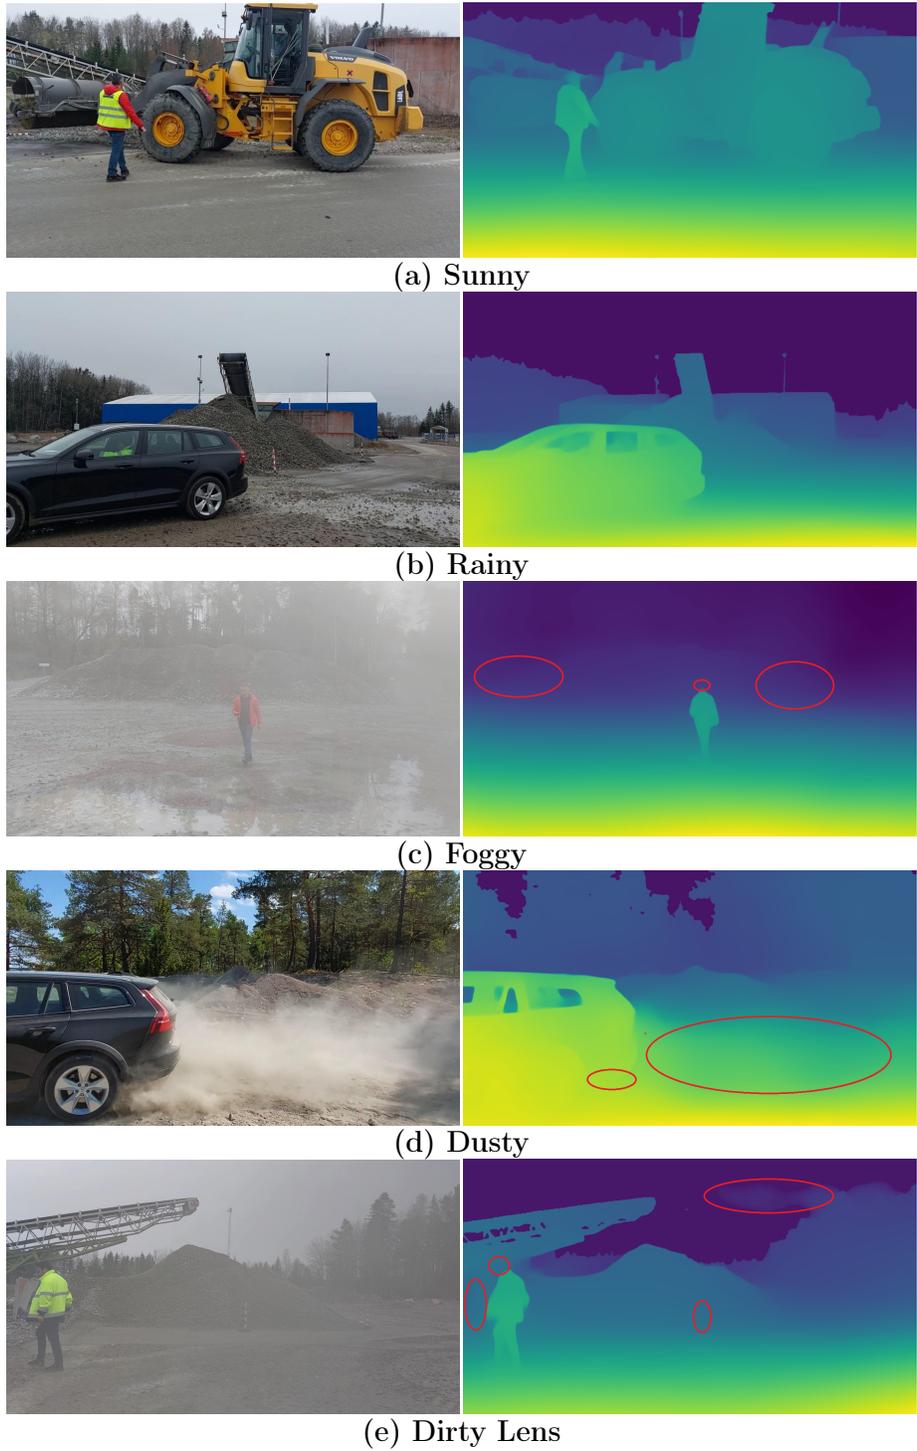


Fig. 2: Visual representation of samples from our dataset under various weather conditions, accompanied by the corresponding disparity map. The red circles indicate the disparity region with low-confidence matches.

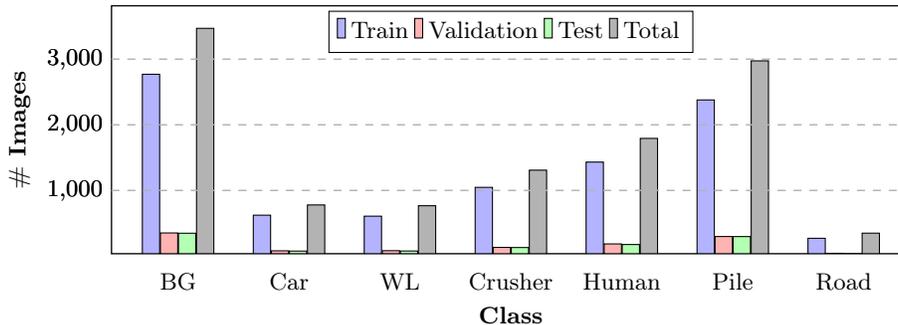


Fig. 3: Distribution of images in each class, including background (BG), Car, wheel-loader (WL), crusher, human, pile, and road.

5 Experimental Results

5.1 Evaluation Metrics

We use mean intersection over union (mIoU) as the evaluation metric for assessing predictions in our experiments. The Mean Intersection over Union (mIoU) used in semantic segmentation is defined as $mIoU = \frac{1}{N} \sum_{i=1}^N \frac{|X_i \cap Y_i|}{|X_i \cup Y_i|}$, where N , X , and Y are the total number of classes, the set of pixels predicted to belong to class i , and the set of pixels in the ground truth that belong to class i , respectively.

5.2 Experimental Setup

Training Dataset. We conduct evaluations on two versions of ConstScene: (i) $D1$, representing the original dataset without any data augmentation; and (ii) $D2$, denoting the original dataset incorporating the first three augmentation techniques (*blurring + pepper-and-salt noise + regions removal*).

Model Details. For evaluation, we leverage U-Net model [21] with ResNet18 and ResNet50 [11] encoder architectures; and SegFormer [29] as a state-of-the-art Vision Transformer (ViT) [7] model. A U-Net is a CNN architecture designed for image segmentation tasks. It features a U-shaped structure, comprising a contracting path for capturing high-level features and a corresponding expansive path for precise segmentation. Skip connections maintain spatial details. The U-Net effectively balances local and global contexts, making it versatile for various computer vision applications. In our exploration of ViT, we examined various iterations of the advanced SegFormer model [29]. SegFormer has two main modules: a hierarchical transformer encoder; and a lightweight multi-layer-perception decoder to predict the final segmentation mask.

Training Details. Table 2 summarizes training parameters for U-Net and SegFormer models. We utilize NVIDIA[®] A5000 for training U-Net and SegFormer models on ConstScene. Considering sustainability, it’s noteworthy to

highlight the carbon emissions associated with the training process. By utilizing the machine learning CO₂ impact calculator tool ³, the total carbon footprint for all experiments is roughly 8.65 kg CO₂.

Table 2: Summary of training parameters for U-Net and SegFormer models.

Parameter	U-Net [21]	SegFormer [29]
Optimizer	Adam	Adam
Initial Learning Rate	3×10^{-4}	2×10^{-5}
# Epochs	50	64
Batch Size	{1000, 1500}	{200, 500}

5.3 Prediction Results

Analyzing Results of U-Net. Table 3 presents the test results of U-Net with various backbone architectures trained on the original **ConstScene** with and without the incorporation of data augmentation techniques. Results show that U-Net with the ResNet50 encoder trained on the non-augmented dataset (*D1*) provides the best results (55.48%) compared to other training settings. In addition, data augmentation (*D2*) provides up to 11.59% higher accuracy for U-Net with the ResNet50 encoder demonstrating that the model excels at capturing the diverse patterns found in real-world.

Table 3: Performance of U-Net [21] with different backbone architectures trained on **ConstScene**.

Encoder	# Parameters	Dataset	Batch Size	Train Time (Hour)	mIoU (%)
ResNet18	13.7M	<i>D1</i>	1500	1	51.93
ResNet18	13.7M	<i>D2</i>	1500	1.7	55.76
ResNet50	31M	<i>D1</i>	1000	2.5	55.48
ResNet50	31M	<i>D2</i>	1000	3	67.07

Analyzing Results of SegFormer. Table 4 shows the test results of SegFormer trained on **ConstScene**. The findings reveal that SegFormer-B5, when trained on the original dataset (*D1*), achieves an accuracy of 82.27%, surpassing SegFormer-B0’s results by an improvement of 35.3%.

Analyzing Results of Data Augmentation. When SegFormer-B5 undergoes training using the dataset enhanced with data augmentation *D2*, it yields superior accuracy compared to the non-augmented dataset. This suggests that

³ <https://mlco2.github.io/impact/>

the *blurring + pepper-and-salt noise + regions removal* augmentation techniques contribute to the model’s improved ability to learn variations present in real-world data.

Table 4: Performance of SegFormer [29] with different model sizes trained on ConstScene.

SegFormer-*	# Parameters	Dataset	Batch Size	Train Time (Hour)	mIoU (%)
B0	3.7M	D1	500	3.7	46.96
B0	3.7M	D2	500	7.5	64.84
B5	84.6M	D1	200	13	82.27
B5	84.6M	D2	200	27	87.09

5.4 Advantages of ConstScene in Defending Against Adversarial Attacks

Adversarial attacks [5, 22, 23] involve the deliberate manipulation of input data to cause a model to produce incorrect or misleading outputs. Adversarial attacks can significantly threaten the reliability and security of machine learning systems. Challenging weather conditions and natural phenomena, such as poor lighting and shadow, are categorized as a type of black-box adversarial attack [27, 32]. ConstScene can be leveraged to improve the model’s robustness against such adversarial attacks through the adversarial training strategy. By deliberately exposing the model to adverse weather samples during training, (i) the model learns to defend against such attacks and becomes more adept at recognizing and mitigating adversarial inputs; and (ii) the model is better equipped to handle variations and distortions in real-world data, making it more resilient against adversarial inputs.

6 Conclusion and Future Work

We presented a novel large-scale dataset, dubbed ConstScene, for studying semantic object detection in construction environment across various weather and environmental conditions with extensive annotations for comprehensive understanding and analysis. We have extended our dataset by employing data augmentation techniques aligned with environmental hazards prevalent in construction sites. We also showed the results of two popular semantic segmentation models trained on our proposed dataset. In future work, semantic segmentation in construction environments can be further improved by (i) considering objects from other construction segments such as mining; and (ii) generating synthetic data using generative adversarial networks. Our new findings and the proposed dataset pose a high potential in stimulating the development of innovative techniques to enhance the efficacy of object detection models against adversarial attacks.

Acknowledgements This work was supported by Volvo Construction Equipment AB, and the Swedish research financier for universities, KK-stiftelsen, through the SACSys project.

References

1. Ahmed, M., Hashmi, K.A., Pagani, A., Liwicki, M., Stricker, D., Afzal, M.Z.: Survey and performance analysis of deep learning based object detection in challenging environments. *Sensors* **21**(15), 5116 (2021)
2. Alomar, K., Aysel, H.I., Cai, X.: Data augmentation in classification and segmentation: A survey and new strategies. *Journal of Imaging* **9**(2), 46 (2023)
3. Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9297–9307 (2019)
4. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11621–11631 (2020)
5. Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., Mukhopadhyay, D.: Adversarial attacks and defences: A survey. *arXiv preprint arXiv:1810.00069* (2018)
6. 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*. pp. 3213–3223 (2016)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
8. Duan, R., Deng, H., Tian, M., Deng, Y., Lin, J.: Soda: A large-scale open site object detection dataset for deep learning in construction. *Automation in Construction* **142**, 104499 (2022)
9. Dwyer, B., Nelson, J., Solawetz, J., et al.: Roboflow (version 1.0). *Software* (2022), available from <https://roboflow.com>
10. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* **32**(11), 1231–1237 (2013)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
12. Huang, X., Cheng, X., Geng, Q., Cao, B., Zhou, D., Wang, P., Lin, Y., Yang, R.: The apolloscape dataset for autonomous driving. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 954–960 (2018)
13. Kim, J., Chi, S.: Robust real-time object detection on construction sites using integral channel features. In: *Proceedings of the 6th International Conference on Construction Engineering and Project Management (ICCEPM 2015)*, Busan, Korea. pp. 304–309 (2015)
14. Kolar, Z., Chen, H., Luo, X.: Transfer learning and deep convolutional neural networks for safety guardrail detection in 2d images. *Automation in Construction* **89**, 58–70 (2018)

15. Loni, M., Zoljodi, A., Maier, D., Majd, A., Daneshtalab, M., Sjödin, M., Juurlink, B., Akbari, R.: Densedisp: Resource-aware disparity map estimation by compressing siamese neural architecture. In: 2020 IEEE Congress on Evolutionary Computation (CEC). pp. 1–8. IEEE (2020)
16. Loni, M., Zoljodi, A., Majd, A., Ahn, B.H., Daneshtalab, M., Sjödin, M., Esmaeilzadeh, H.: Faststereonet: A fast neural architecture search for improving the inference of disparity estimation on resource-limited platforms. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **52**(8), 5222–5234 (2021)
17. Meng, L., Gang, Z., Yawei, Y., Jun, S.: Survey of object detection methods under adverse weather conditions. *Journal of Computer Engineering & Applications* **58**(13) (2022)
18. Ming, Y., Meng, X., Fan, C., Yu, H.: Deep learning for monocular depth estimation: A review. *Neurocomputing* **438**, 14–33 (2021)
19. Nath, N.D., Behzadan, A.H.: Deep convolutional networks for construction object detection under different visual conditions. *Frontiers in Built Environment* **6**, 97 (2020)
20. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* **44**(3), 1623–1637 (2020)
21. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)
22. Salimi, M., Loni, M., Sirjani, M.: Learning activation functions for adversarial attack resilience in cnns. In: *International Conference on Artificial Intelligence and Soft Computing*. pp. 203–214. Springer (2023)
23. Salimi, M., Loni, M., Sirjani, M., Cicchetti, A., Abbaspour Asadollah, S.: Saraf: Searching for adversarial robust activation functions. In: *Proceedings of the 2023 6th International Conference on Machine Vision and Applications*. pp. 174–182 (2023)
24. Somua-Gyimah, G., Frimpong, S., Nyaaba, W., Gbadam, E.: A computer vision system for terrain recognition and object detection tasks in mining and construction environments. In: *SME annual conference* (2019)
25. Sun, P., Kretschmar, H., Dotiwala, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2446–2454 (2020)
26. Tajeen, H., Zhu, Z.: Image dataset development for measuring construction equipment recognition performance. *Automation in Construction* **48**, 1–10 (2014)
27. Wang, D., Yao, W., Jiang, T., Tang, G., Chen, X.: A survey on physical adversarial attack in computer vision. *arXiv preprint arXiv:2209.14262* (2022)
28. Xiao, B., Kang, S.C.: Development of an image data set of construction machines for deep learning object detection. *Journal of Computing in Civil Engineering* **35**(2), 05020005 (2021)
29. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems* **34**, 12077–12090 (2021)
30. Xuehui, A., Li, Z., Zuguang, L., Chengzhi, W., Pengfei, L., Zhiwei, L.: Dataset and benchmark for detecting moving objects in construction sites. *Automation in Construction* **122**, 103482 (2021)

31. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2636–2645 (2020)
32. Zhong, Y., Liu, X., Zhai, D., Jiang, J., Ji, X.: Shadows can be dangerous: Stealthy and effective physical-world adversarial attack by natural phenomenon. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15345–15354 (2022)