Safe Traffic Sign Recognition through Data Augmentation for Autonomous Vehicles Software

Lisa Jöckel
Fraunhofer Institute for Experimental
Software Engineering IESE
Germany
lisa.joeckel@iese.fraunhofer.de

Michael Kläs
Fraunhofer Institute for Experimental
Software Engineering IESE
Germany
michael.klaes@iese.fraunhofer.de

Silverio Martínez-Fernández
Fraunhofer Institute for Experimental
Software Engineering IESE
Germany
silverio.martinez@iese.fraunhofer.de

Abstract—**Context: Since autonomous vehicles operate in an open context, their software components, including data-driven ones, have to reliably process inputs (e.g., obtained by cameras) in order to make safe decisions. A key challenge when providing reliable data-driven components is insufficient training data, which could lead to wrong interpretation of the environment, thereby causing accidents. Aim: The goal of our research is to extend available training data of data-driven components for safe autonomous vehicles using the example of traffic sign recognition. Method: We developed an approach to create realistic image augmentations of various quality deficits and applied them on the German traffic sign recognition benchmark dataset (GTSRB). Results: The approach results in images augmented with (any combination of) seven different quality deficits affecting traffic sign recognition (rain, dirt on lens, steam on lens, darkness, motion blur, dirt on sign, backlight) and considers dependencies between combined quality deficits and influences from other contextual information. Conclusion: Our approach can be used to obtain more comprehensive datasets, especially also including samples with quality deficits that are difficult to gather. By structuring the augmentation into a set of basic components, the approach can be adapted for other application domains (e.g., person detection).**

Keywords—safety, autonomous driving, image augmentation, machine learning, convolutional neural networks, data quality

I. INTRODUCTION

Autonomous vehicles have to operate in an open context, so they process sensor data, especially from camera sensors, to perceive their environment and collaborate with other traffic participants. Data-driven components including, e.g., convolutional neural networks, are emerging in automotive software engineering [1]. In contrast to traditionally engineered components, the functional behavior of these components is derived from training data, which introduces several sources of uncertainty [2]. Therefore, training data should (a) reflect the intended application scope [3] and (b) be comprehensive with respect to relevant quality deficits [4].

In traffic sign recognition, possible quality deficits caused by context conditions, the sign itself, or the camera sensor were identified in [5]. Examples are weather conditions, physical damage, dirt on the lens, or lighting conditions. Image augmentations for street scenes done so far are either limited to a few quality deficits [6] or lack realism [7]. The possible occurrence of multiple quality deficits together is considered to some extent in [8], where different augmentations are combined using the LAB color space.

In this context, we propose an approach for creating methods for augmentation that can be applied in combination, and whose components can be adapted and reused for further quality deficits. The proposed approach can be utilized to provide augmentation functionality in frameworks like [5].

II. APPROACH: QUALITY DEFICIT AUGMENTATION

This section describes the five steps of quality deficit augmentation for traffic sign images in the context of traffic sign recognition: understanding of the manifestation, determination of appearance parameters, layer creation, layer consolidation, and deficit combination. The images we use to illustrate the approach are taken from the German traffic sign recognition benchmark dataset (GTSRB) [9].

A. Quality Deficit Manifestation

In order to change an image to augment a certain quality deficit, we have to understand how the deficit manifests in an image [5]. Two different approaches are illustrated in the following. First, rain appears as streaks on an image (cf. Fig. 1). For real rain, those streaks are not identical. There are small variations in the length and direction of the streak lines. Besides, opacity varies depending on the distance from the camera. A dirty camera lens creates a far more irregular structure. To get a naturally looking structure, one possibility is to incorporate simplex noise (cf. Fig. 1). To increase realism and avoid that the model learns fixed structures, simplex noise and rain streaks are created individually for every image. In both cases, we add new content to the image. Second, a different approach is utilized for motion blur. A shrunken version of the image is blended with the original image to get a zoom-like effect (cf. Fig. 1). Hence, no new content is added, but the image itself is used in a transformed version to change the original image.

Some general techniques can be adjusted and reused for various kinds of quality deficits. Simplex noise, e.g., can be adapted to provide naturally looking structures or stochastically add irregularities to quality-deficient representations that otherwise would look artificial. It can be integrated into the augmentation of steam on the camera lens, e.g., which would otherwise be depicted as a monotonous gray veil (cf. Fig. 1). Image brightness and saturation can be easily manipulated in the HSV color space by increasing or decreasing the value and saturation components.
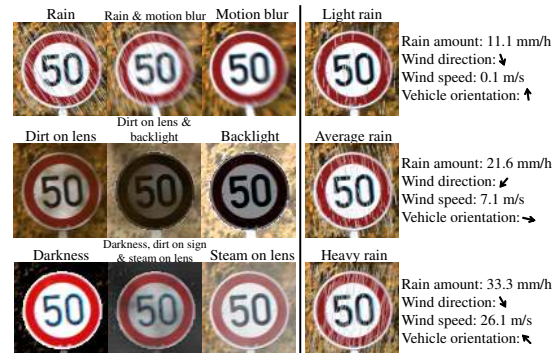

				Rain amount: 11.1 mm/h Wind direction: ↘ Wind speed: 0.1 m/s Vehicle orientation: ↑
Rain	Rain & motion blur	Motion blur	Light rain	
Dirt on lens	Dirt on lens & backlight	Backlight	Average rain	Rain amount: 21.6 mm/h Wind direction: ↙ Wind speed: 7.1 m/s Vehicle orientation: →
Darkness	Darkness, dirt on sign & steam on lens	Steam on lens	Heavy rain	Rain amount: 33.3 mm/h Wind direction: ↓ Wind speed: 26.1 m/s Vehicle orientation: ↙

Fig. 1. Example traffic sign with augmentations for quality deficit combinations, and rain in different contexts.

B. Parameter Determination

The appearance of a quality deficit on a concrete sign is also influenced by the following three aspects:

a) Influence of Quality Deficit Intensity: Regarding the quality deficit rain, the intensity influences the number of rain streaks, and consequently also the number of lines visible in the image, as well as the length and thickness of the lines. For a dirty or steamed-up lens, the intensity affects the general opacity of the dirt or steam or the proportion of the image area to which the augmentation is applied (e.g., partially dirty).

b) Influence of Scope Characteristics: Besides intensity, characteristics of the specific application situation can also influence other aspects of the appearance of a quality deficit. Wind, e.g., causes a slant depending on the direction of the driving and the wind, and the type of environment (e.g., forest or city) affects the color of the dirt on a sign [5].

c) Influence of Quality Deficit Combination: The presence of one quality deficit can change the appearance of another. When it is dark, steam on a lens will also appear darker than otherwise. Thus, an application order of quality deficit augmentations and the effects between multiple quality deficits present have to be considered [5].

C. Opacity Layer

Opacity is high where effects from the quality deficit are visible in the image. In the opacity layer of the quality deficit dirt on sign (cf. Fig. 2), pixels with high opacity (white or light gray) correspond to the dirt's location, while pixels with low opacity (black or dark gray) are hardly affected. Some quality deficits only occur on the sign itself. In the GTSRB dataset, traffic signs are already cut by their bounding box, hence the traffic sign occupies most of the image. We used the OpenCV GrabCut algorithm [10] to get a pixel mask of the traffic sign. Known foreground and background image parts are selected depending on the geometric shape, which is derived from the sign type. The image of a sign or lens covered only partially with dirt, e.g., can be augmented by blending a gradient into the opacity layer (like a linear gradient from zero to one).

D. Color Layer

While the opacity layer defines which image pixels are affected by the quality deficit, the color layer defines in which way they are affected. This can be a single color for the whole image, as for dirt (cf. Fig. 2), but a distinction between traffic sign and background is also possible. Backlight manifests in the image with a brighter background, but darker foreground. As a safety feature, signs are coated with reflective foil to ensure perceptibility at night [11]. Hence, a reflection mask is introduced for the quality deficit darkness. Initially, it is set to the pixel mask of the traffic sign and is adjusted if the reflection property is affected, e.g., by dirt on the sign.

E. Layer Consolidation and Quality Deficit Combination

To consolidate opacity and color information, color layer and image are blended together with opacity values as weights for each pixel. In the case of separate foreground and background color layers, the same is applied to pixels of the

traffic sign and of the background, respectively. As multiple quality deficits can occur together, their augmentations cannot be considered to be independent, but rather build upon each other. Consequently, information from one quality deficit augmentation has to be passed on to the next, including the augmented image and the reflection mask. Each quality deficit augmentation also gets the pixel mask of the traffic sign, the original image, and information about other quality deficits.

III. CONCLUSION AND ROADMAP

We presented realistic augmentations for seven quality deficits. They can be combined with each other considering dependencies between them, and incorporate contextual information. The results of using GrabCut for traffic sign segmentation considering the sign geometry are promising.

Future work will target several areas. Enhancements can be achieved by including the distance between camera and traffic sign. Depth information is not available, but so far, the size of the bounding box around the traffic sign is used as an approximation to determine rain streak length. Regarding the augmentation of a physically damaged sign (e.g., broken or bent), there is the need to deal with deleting parts of the sign and replacing them with realistic background. Another difficulty arises for inserting objects that partly occlude the traffic sign, e.g., a tree branch.

REFERENCES

[1] M. Borg, C. Englund, K. Wnuk, B. Duran and C. Levandowski, "Safely entering the deep: a review of verification and validation for machine learning and a challenge elicitation in the automotive industry," Journal of Automotive Software Engineering, vol. 1, no. 1, pp. 1-19, 2019.

[2] M. Kläs and A. M. Vollmer, "Uncertainty in Machine Learning Applications – A Practice-Driven Classification of Uncertainty," in WAISE 2018, pp. 431-438, 2018.

[3] M. Kläs and L. Sembach, "Uncertainty wrappers for data-driven models," in WAISE 2019, in press.

[4] Q. Rao and J. Frtunikj, "Deep learning for self-driving cars: chances and challenges," in Proceedings of the 1st International Workshop on Software Engineering for AI in Autonomous Systems, pp. 35-38, 2018.

[5] L. Jöckel and M. Kläs, "Increasing trust in data-driven model validation," in SafeComp 2019, in press.

[6] Z. Pezzementi, T. Tabor, S. Yim, J. K. Chang and B. Drozd, "Putting image manipulations in context: robustness testing for safe perception," in IEEE SSRR, 2018.

[7] C.-H. Cheng, C.-H. Huang and G. Nührenberg, nn-dependability-kit: engineering neural networks for safety-critical systems, https://arxiv.org/abs/1811.06746, 2018.

[8] S. P. Harisubramanyabalaji, S. ur Réhman, M. Nyberg and J. Gustavsson, "Improving image classification robustness using predictive data augmentation," in SAFECOMP 2018, pp. 548-561, 2018.

[9] "German Traffic Sign Benchmarks," [Online]. Available: http://benchmark.ini.rub.de/?section=gtsrb. [Accessed 9 Juni 2019].

[10] "OpenCV: Interactive Foreground Extraction using GrabCut Algorithm," [Online]. Available: https://docs.opencv.org/3.4.3/d8/d83/tutorial_py_grabcut.html. [Accessed 9 Juni 2019].

[11] Ł. Jeliński, J. Wachnicka, J. Jamroz, M. Kalisz and P. Kaźmierczak, "Testing the durability and function of road traffic management devices," in MATEC, 2017.

No augmentation   Pixel mask (sign)   Reflection mask   Opacity layer   Color layer   Augmented image

Fig. 2. Example traffic sign with dirt on sign augmentation, segmented traffic sign, reflection mask after augmentation, opacity, and color layer.