

Detection of Active Emergency Vehicles using Per-Frame CNNs and Output Smoothing

Meng Fan, Craig Bidstrup, Zhaoen Su, Jason Owens, Gary Yang, Nemanja Djuric
Aurora Innovation, Inc.

{meng.fan, cbidstrup, zhaoen, jlowens, garyyang, ndjuric}@aurora.tech

Abstract

While inferring common actor states (such as position or velocity) is an important and well-explored task of the perception system aboard a self-driving vehicle (SDV), it may not always provide sufficient information to the SDV. This is especially true in the case of active emergency vehicles (EVs), where light-based signals also need to be captured to provide a full context. We consider this problem and propose a sequential methodology for the detection of active EVs, using an off-the-shelf CNN model operating at a frame level and a downstream smoother that accounts for the temporal aspect of flashing EV lights. We also explore model improvements through data augmentation and training with additional hard samples.

1. Introduction

Self-driving vehicles (SDVs) rely on an onboard perception system to understand their surroundings, which includes detection of all traffic actors in the vicinity as well as their tracking through time. The output of this system commonly includes actor type (such as vehicle, pedestrian, bicyclist, etc.), as well as the state of the actor (such as bounding box size, position, velocity, acceleration, etc.), which is a well-studied problem in the literature [7,9]. However, these outputs may not capture all relevant attributes for some traffic actors that the SDV may interact with. This is the case when it comes to the detection of active emergency vehicles (EVs), where the state of the emergency lights provides an important additional piece of information to the SDV.

Vision-based detection of EVs has been a topic of several recent studies. Most researchers proposed to first use an object detector implemented as a convolutional neural network (CNN) that is then run on an input camera image to detect vehicle actors, followed by an EV-specific classifier operating on cropped images of detected actors. In particular, the authors of [19] proposed to use YOLO [18] to detect vehicles and crop images around them, and then

apply a VGG-16-based EV classifier trained on individual vehicle images. In [4] the authors introduced a similar two-stage approach with an object detector and a CNN classifier, which is trained and tested on a small custom data set, while [22] used a single-stage model by training an object detector YOLOv4 [5] directly on a data set of two object classes: EV and non-EV. In [10,13] the authors explored several existing CNN architectures (such as DenseNet [12]) for EV classifications using image crops of individual vehicles as inputs, trained and evaluated on smaller data. In addition, the authors of [10] showed the benefits of using pre-trained weights obtained by training on ImageNet data [8], followed by a fine-tuning step used to specialize the models for EV-specific tasks. The authors in [17] extracted color values from the top of an image crop (where the EV siren lights are expected to be found) provided by a vehicle detector and applied a support vector machine for classification.

In these earlier works the data came either from non-EV-specific public sources (such as MS COCO [15]) or from EV-specific YouTube channels [1,2,20] and CCTV footage. The statistics in such data sets are far from what we find in the real world though. For example, the EV-to-non-EV ratio is unrealistically high in data sets collected from EV-specific YouTube channels, and in general the data will not be able to cover the numerous non-EV types and appearances commonly observed on the roads. Moreover, such data are usually biased towards close-range EVs that are easier to classify, while for the SDV use-case one would like to detect EVs at the first sight at long ranges, which presents a more challenging task. Besides, the previous approaches did not consider the EV activeness (i.e., whether or not the EV lights are flashing). In this work we adopt the two-step approach discussed above while addressing several limitations of the prior studies. We collect on-road data from our self-driving fleet covering rich scenes involving EVs, including day and night, short and long ranges, highway and urban, and various other traffic conditions. We focus on detecting active EVs that are more important for SDVs to respond to. In addition, we apply approaches such as data augmentation, hard-sample mining, and output smoothing,

Table 1. Data set summary

Category	Distribution
Time of logs	Day: 81.1%, Night: 18.9%
Vehicle type	EV: 3.4%, non-EV: 96.6%
EV type	police vehicle: 80.0%, fire: 13.4%, ambulance: 6.6%
EV activeness	active: 90.0%, inactive: 10.0%
Bulb state of active EVs	bulb-on: 91.8%, bulb-off: 8.2%

which help further boost the performance.

An alternative approach worth mentioning is audio-based detection [3, 6, 21, 22], which does not require an unobstructed line of sight to detect active EVs. However, on highways and in general at higher ego-speeds acoustic approaches are significantly affected by ambient noise (such as wind). We do not consider such methods in our current work and instead focus on image-based approaches.

2. Data set and labeling

We collected on-road data logs using an autonomous driving fleet equipped with lidar and camera sensors. For all vehicles we perform 4D labeling (i.e., a sequence of 3D bounding boxes) and 2D labeling (i.e., a 2D bounding box within a camera image). We associate 4D and 2D labels and generate a collage video of each vehicle using the camera images. We then use it to label the actor into one of 4 categories: police car, ambulance, fire truck, and non-EV. If the vehicle is labeled as an EV it will trigger the second-stage labeling for activeness and bulb state. In particular, at each time frame we label the vehicle as active if the beacon light is flashing, and inactive otherwise. For active frames, we label it as "bulb on" if the light is illuminated and "bulb off" otherwise. Note that the activeness labeling relies on the local temporal context, while the bulb state labeling only relies on a single frame.

We collected logs containing a total of around 12,000 unique vehicles, each with a maximum duration of 25 seconds. As seen in the data summary shown in Table 1, we observed a large fraction of non-EVs, which is more realistic than previous studies. We also find that the majority of EVs are active. Active EVs usually have multiple beacon bulbs flashing, and thus when looking at each single frame there is a small chance (8.2%) to observe all bulbs being off. This motivates us to design a single-frame detector for active EV actors, which also has faster inference than a multi-frame detector.

3. Methodology

3.1. The EV detection system

The EV detection system is composed of three phases: pre-processing, inference, and post-processing. During pre-processing, a tracker system produces tracks of the surrounding vehicles that are used as input. We project the 3D coordinates of the corners of each vehicle track’s bounding box into the camera image and compute the centroid. For each centroid within the image bounds, we use the smallest axis-aligned square that encapsulates the projected corners to crop the image and resize it to 224×224 pixels. We ignore tracks where the centroid is out of the field of view of the forward camera, as well as those with a projected width smaller than a certain threshold (set to 18 pixels in this study). We then batch the valid image patches and forward them to the inference module. During inference, we run the EV classifier model (discussed in detail in the next section) on the image patches and output probabilities indicating whether a certain vehicle is an active EV or not. Then, in order to capture the temporal component of active flashing lights, we implemented a downstream smoother in the post-processing phase that infers the final overall state. The smoother is a cyclic buffer that keeps a ledger of the last 25 valid EV outputs for each vehicle track’s history and outputs a smoothed result. We require that there are at least 6 detected frames of the actor and more than $T = 50\%$ of per-frame active EV outputs in the buffer to mark the vehicle as being an active EV. This helps suppress transient positive outputs and thus mitigate false positives (FPs), discussed further in Section 4. The EV detection system can handle hundreds of actors (processed as a batch) in real-time, with an average latency of less than 10 milliseconds.

3.2. The EV classifier model

The EV classifier is a CNN using ResNet-18 (layers 1 to 4) [11] as a backbone to extract image features, followed by a global average pooling layer and a fully-connected layer with a single output channel that is passed through sigmoid to output probability. The ResNet backbone uses the weights pre-trained on the ImageNet data [8] and is not frozen during training. We set the label as positive when a vehicle is an EV with a "bulb on" state at the time of image capture, and as negative otherwise. We use focal loss [14] due to a large data imbalance. The Adam optimizer with 0 weight decay and an initial learning rate of 10^{-4} is used. We apply a customized optimizer wrapper that decreases the learning rate when there is no loss improvement for a certain number of iterations, and stops training when the learning rate drops to a given low value. The collected logs are split randomly into train and test sets with a 3 : 1 ratio. The model is implemented in PyTorch [16] and trained on a single GPU.

Table 2. Per-frame classifier results (A: data augmentation; M: mined data from data engine)

Model	% change of max-F1	% change of precision at 0.8 recall
Baseline	0	0
+ A	3.11	5.3
+ A + M	5.73	10.4

3.3. Techniques for data improvement

The considered problem is non-trivial due to the imbalance between EVs and non-EVs and challenging image artifacts, especially at long ranges. We applied several techniques to address these issues.

3.3.1 Data augmentation

Data augmentation is often applied when training an image classifier by modifying raw images using operations such as cropping, flipping, rotation, resizing, and others. We augment only the rare active-EV class and focus on approaches that do not introduce unrealistic deviations and additional bias between the classes. Our solution updates the vehicle track bounding box as follows. The track state contains x , y , and z coordinates of the bounding box center, as well as its length, width, and height. To each of these states we fit a normal distribution on the training set. Then, given an input track, we sample from this distribution to generate new bounding boxes, ensuring the augmented data is likely to be seen in the real world. When creating the final train data set we use a sampling ratio of $2\times$ on positive examples, and also downsample negative examples by $5\times$.

3.3.2 Data engine: Mine, label, and retrain

As EVs are very rare compared to other traffic actors, we explored additional sourcing methods to find relevant data in the collected data logs. Another practical problem is that the proposed model may produce FPs, and we would like to collect such hard samples to help train a better classifier. To address these two problems, we developed a data engine to continuously improve the learned model. In particular, we mine the data in the newly collected logs to find the events with active-EV detections by the EV classifier, regardless of whether or not they were true or false positives. Then, we label such events and add them to the data set, before retraining the EV classifier model.

4. Results

We performed hyper-parameter tuning to train a strong initial model, which is used as a baseline to which other



Figure 1. Example detections (top row: true positives, bottom row: difficult true negatives)

Table 3. Per-actor results of the output smoother for various smoother threshold T

T	% change of precision	% change of recall	% change of F1 score
0%	0	0	0
30%	10.30	-0.68	5.07
50%	9.87	-0.51	4.87
70%	6.80	-20.0	-7.02

experiments are compared. The results are shown in Table 2. By using data augmentation discussed in Section 3.3.1, we see 3.11% boost in the max-F1 score compared to the baseline. If we also add more training data collected through hard-sample mining (nearly doubling the number of vehicles) and retrain the model, we see an improvement of 5.73% in max-F1. Moreover, this method improves the precision by more than 10% at a fixed 0.8 recall. More importantly, the data engine helps to continuously iterate and improve the model by adding useful data as it is collected.

In Figure 1 we show examples of true positives (TPs) and difficult true negatives (TNs) of the improved "+ A + M" model from Table 2. The TPs that are shown in the first row cover different types of EVs, various weather and light conditions, as well as multiple viewing angles. The difficult negative cases (shown in the second row) were false positives of the baseline model, but then successfully predicted as TNs by the improved model. It is interesting to note that such difficult cases were usually caused by some other light sources, including brake lights from vehicles queued ahead, flashing lights of construction vehicles, or sun reflections.

Beyond the learned EV model, we also evaluated the performance of the downstream smoother. Using a different setup of the smoother we show the corresponding changes in precision, recall, and F1 score in Table 3, where we tune the threshold T (i.e., the minimum fraction of positive outputs in the buffer required to output a final positive state). As can be seen, when the threshold is 30% and 50%, we reach a higher F1 score and precision with very small regression in recall, which shows the effectiveness of smoother in mitigating FPs. However, when the threshold

is set to a high value (such as 70%), the system starts to generate more FNs which hurts the overall performance.

5. Conclusion

We considered the problem of detecting active emergency vehicles in the context of self-driving vehicles. We proposed to use a frame-level EV detector whose outputs are fed to an output smoother, which captures the temporal dimension of such actors. Finally, we evaluated the method on large-scale, real-world data, and made further improvements by data augmentation and by training with additional hard samples mined using a data engine.

References

- [1] *Youtube Channel: TGG—Global Emergency Responses. [Online]. Available: <https://www.youtube.com/c/TGGGlobalEmergencyResponses>, 2020.* 1
- [2] *Youtube Channel: Demonracer2. [Online]. Available: <https://www.youtube.com/user/demonracer2>, 2020.* 1
- [3] Muhammad Asif, Muhammad Usaid, Munaf Rashid, Tabarka Rajab, Samreen Hussain, and Sarwar Wasi. Large-scale audio dataset for emergency vehicle sirens and road noises. *Scientific data*, 9(1):1–9, 2022. 2
- [4] Abhishek Baghel, Aprajita Srivastava, Aayushi Tyagi, Somya Goel, and Preeti Nagrath. Analysis of ex-yolo algorithm with other real-time algorithms for emergency vehicle detection. In *Proceedings of First International Conference on Computing, Communications, and Cyber-Security (IC4S 2019)*, pages 607–618. Springer, 2020. 1
- [5] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020. 1
- [6] Michela Cantarini, Leonardo Gabrielli, and Stefano Squartini. Few-shot emergency siren detection. *Sensors*, 22(12):4338, 2022. 2
- [7] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 1
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1, 2
- [9] Duarte Fernandes, António Silva, Rafael Névoa, Cláudia Simões, Dibet Gonzalez, Miguel Guevara, Paulo Novais, João Monteiro, and Pedro Melo-Pinto. Point-cloud based 3d object detection and classification methods for self-driving applications: A survey and taxonomy. *Information Fusion*, 68:161–191, 2021. 1
- [10] Samiul Haque, Shayla Sharmin, and Kaushik Deb. Emergency vehicle detection using deep convolutional neural network. In *Proceedings of International Joint Conference on Advances in Computational Intelligence*, pages 535–547. Springer, 2022. 1
- [11] 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. 2
- [12] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 1
- [13] Amine Kherraki and Rajae El Ouazzani. Deep convolutional neural networks architecture for an efficient emergency vehicle classification in real-time traffic monitoring. *IAES International Journal of Artificial Intelligence*, 11(1):110, 2022. 1
- [14] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 2
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1
- [16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 2
- [17] Husniza Razalli, Rusyaizila Ramli, and Mohammed Hazim Alkawaz. Emergency vehicle recognition and classification method using hsv color segmentation. In *2020 16th IEEE International Colloquium on Signal Processing & Its Applications (CSPA)*, pages 284–289. IEEE, 2020. 1
- [18] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 1
- [19] Shuvendu Roy and Md Sakif Rahman. Emergency vehicle detection on heavy traffic road from cctv footage using deep convolutional neural network. In *2019 international conference on electrical, computer and communication engineering (ECCE)*, pages 1–6. IEEE, 2019. 1
- [20] D. Steinhart. *Youtube Channel: World-wide Emergency Responses. [Online]. Available: <https://www.youtube.com/user/wwwrescue911de>, 2020.* 1
- [21] Van-Thuan Tran and Wei-Ho Tsai. Acoustic-based emergency vehicle detection using convolutional neural networks. *IEEE Access*, 8:75702–75713, 2020. 2
- [22] Van-Thuan Tran and Wei-Ho Tsai. Audio-vision emergency vehicle detection. *IEEE Sensors Journal*, 21(24):27905–27917, 2021. 1, 2