Improving Traffic Signal Classification in Safety-Critical Scenarios with Synthetic Data

BACKGROUND

Crashes often occur at intersections because these are the locations where two or more roads cross each other. Activities such as turning left, crossing over, and turning right have the potential for conflicts resulting in crashes. Specifically, the US National Highway Traffic Safety Association (NHTSA) found that **turning left is a leading critical pre-crash event** and occurs in **22.2 percent of all crashes**. Further, approximately 61 percent of crashes that take place while turning or crossing an intersection involve a left-hand turn.[1]

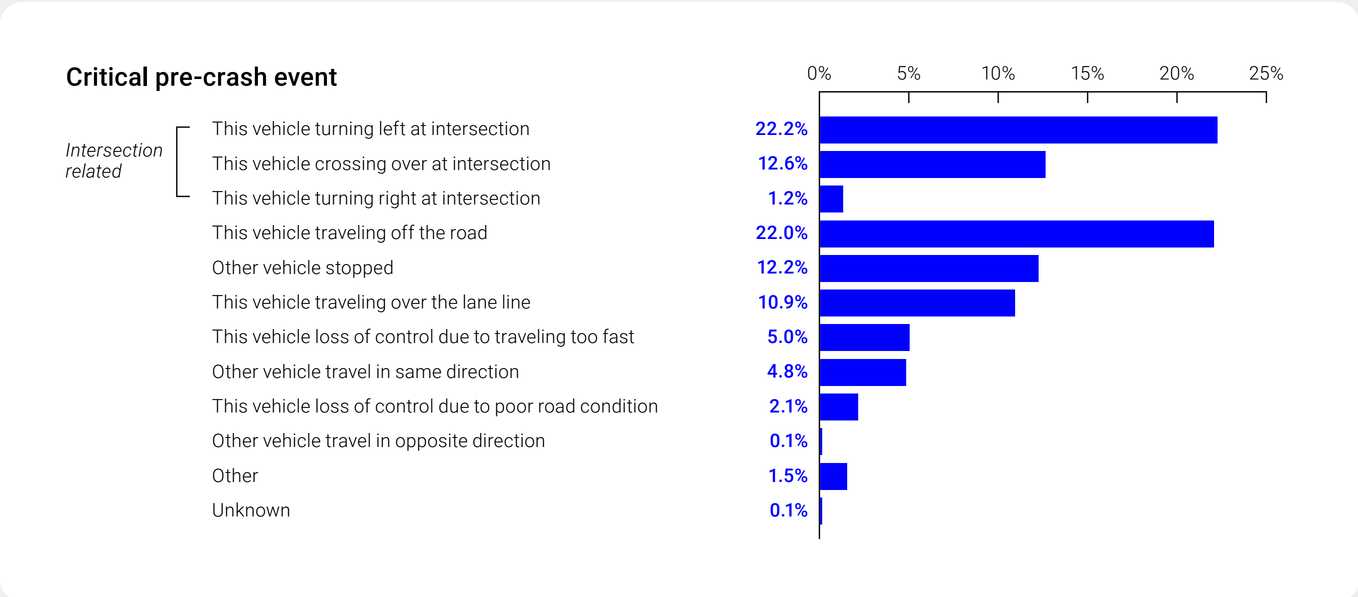

Source : U.S. Department of Transportation - National Highway Traffic Safety Administration

A common approach to increasing safety in left-turn events is the implementation of a **protected left turn**, where a traffic signal indication gives those vehicles intending to make left turns the right to enter the intersection free from conflict with drivers or pedestrians. **Left** arrow signal phases are used.

As Advanced Driver Assistance Systems (ADAS) and Automated Driving Systems (ADS) continue to increase in their abilities to automate driving tasks, the need to train these systems recognize and maneuver complex intersections presents an incredibly difficult challenge. In the case of a protected left turn, this training starts with ensuring that the perception system's traffic signal classifier correctly detects the relevant signal phases. Unfortunately, machine learning models don't always have enough training data because cars don't experience these critical signal phases in the real world as often as they do others, like solid red/yellow/green lights or even red arrows.

Example of a protected left traffic signal

This study examines the impact of class rebalancing by means of synthetic data augmentation, on the classification of protected-left signals by a Faster RCNN Traffic Signal Classifier. In other words, we generated synthetic images of these less-common traffic signal phases and added them to training data collected from the real world.

WHAT ARE TRAFFIC SIGNAL CLASSIFIERS?

Traffic Signal Classifiers are an essential component of all Autonomous Vehicles and many ADAS platforms today. These camera-based systems perform three primary tasks - they identify whether a visible traffic signal exists in the frame, they locate where the signal is in the frame and whether it is facing the vehicle, and they determine which symbols and colors the signal is displaying.

Traffic Signal Classifiers are not new, and many work well in computer vision friendly environments and for providing basic information such as no traffic light, red traffic light, and green traffic light. However, as automated driving systems continue to shift more decision-making and automation from the human driver to the vehicle systems in scenarios like left-hand maneuvers in intersections, the systems' performance requirements increase dramatically.

CHALLENGES IN HIGH-PERFORMANCE TRAFFIC SIGNAL CLASSIFICATION

High-Performance Traffic Signal classifiers must do much more than determine whether a visible signal is on/off or green/red. They must also classify each of the possible colors and the active signal's symbol - is it a solid red or a left-arrow red? The classifiers must also accurately localize the signals - is the traffic signal facing the vehicle? Is it on the near or far side of the intersection? Finally, the Classifier needs to work in real time.

Traffic Signal Classifiers are based on Convolutional Neural Networks, a popular method used in image recognition with deep neural networks (DNNs). As such, training a Traffic Signal Classifier is relatively easy - as long as sufficient training data exists. But therein lies the challenge. Traffic Signals vary tremendously in their densities, quantities, positions, shapes, sizes, orientations, angles, layouts, colors, symbols, maintenance conditions, geographic regions, and more. Collecting enough data to thoroughly train the classifier for all of these variants is undoubtedly costly, time-consuming, and burdensome. In some cases, it can be impossible.

Less Challenging

indirect sunlight without occlusion

More Challenging

harsh sunlight and occlusion

TRADITIONAL APPROACHES TO TRAINING

Today, most Traffic Signal Classifiers are trained on real-world data from commercial or publicly available datasets such as DTLD, BSTLD, LISA, LaRA, or from proprietary data collected by test fleets. The fact that it's expensive to create annotated real-world data means that there's simply less *of* it. The DTLD dataset, for example, comprises over 60,000 instances of traffic signals. Of those, 271 are labeled as yellow left arrows and only 142 are labeled as red-yellow left arrows. 29,547 are labeled as green solid, for comparison.

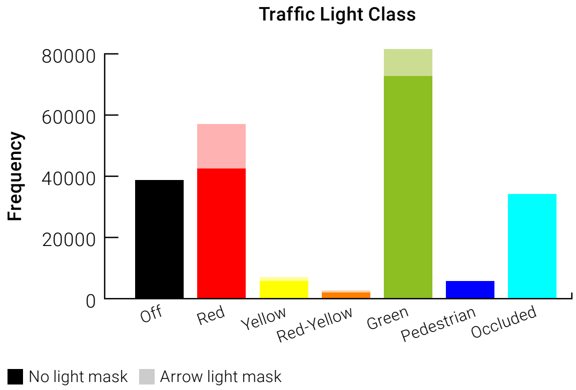

DTLD Signal Phase Label Distribution [3]

Further, while any one set of real-world data has its unique strengths and weaknesses, real-world data generally is inherently limited by its lack of metadata, unbalanced class distributions, incorrect labels, and inaccurate bounding box dimensions. Training AI requires large amounts of accurate data, and errors in the data make it difficult for the models to generalize.

IMPROVING PERFORMANCE WITH SYNTHETIC TRAINING DATA ENRICHMENT

Parameter-driven generation of accurate **synthetic data** creates a previously-impossible solution to the unbalanced distribution problem of annotated real data. This highly-realistic training data is generated on demand, with perfectly-accurate annotations, and fed directly into the same pipeline used for ingesting real-world data.

This allows for training on not only more data, but on more *meaningful* data, that has been parameterized for diversity across not only traffic signal styles, layouts, orientations, and illuminators, but also placement considering shifts, distances, and facing-angles, and environmental conditions.

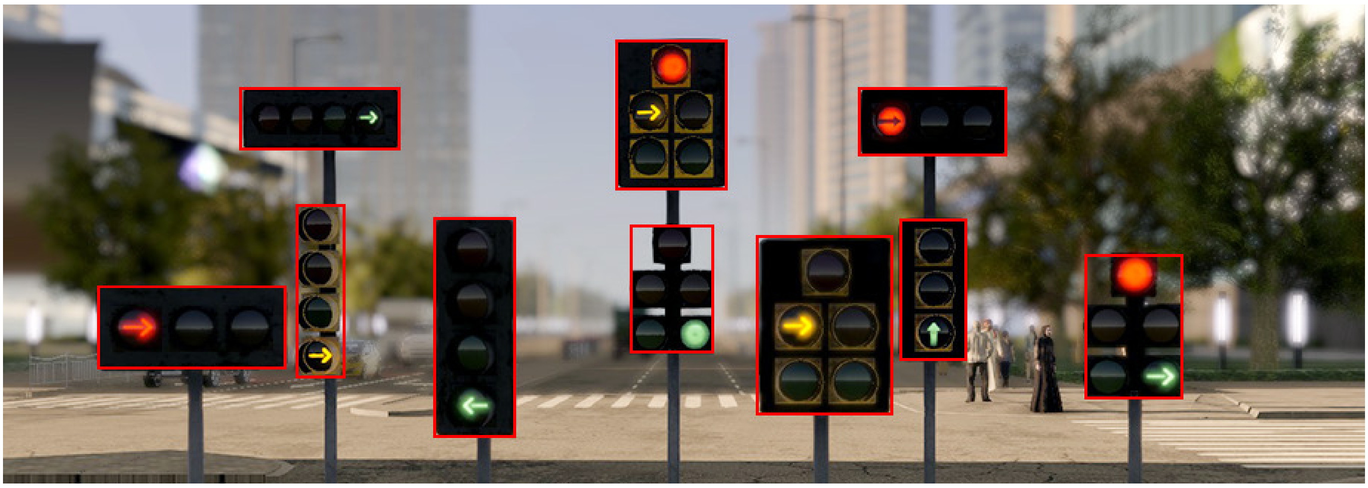

Cognata synthetic traffic signal training data

OUR EXPERIMENT

In this experiment we trained two msCOCO-initialised Faster RCNN networks.

We began by splitting the DTLD dataset into a training set (80% of the instances) and a testing set (20% of the instances). We then trained our network on only the DTLD training set and measured the classifier's Average Precision (AP) across signal phases. AP defines the area under the Precision-Recall curve.

We then generated a small batch of Cognata synthetic data and augmented the DTLD training set to bridge gaps in its coverage of two critical signal phases, yellow left arrows and red-yellow left arrows. Afterward, we trained the network on this enriched DTLD + Cognata Synthetic dataset and again tested the classifier's performance.

Class	Yellow Left Arrow	Red-Yellow Left Arrow
Number of Real Instances (DTLD)	244	115
Number of Synthetic Instances (Cognata)	4265	3996
Ratio of Synthetic to Real Instances	4.18	3.81

Testing performed on DTLD image set

FINDINGS

Enriching the DTLD dataset with only a small amount of synthetic data yielded immediate and noticeable results.

The network trained on DTLD + Cognata significantly outperformed the network trained only on DTLD. The yellow left arrow and red-yellow left arrow classes showed improvements of 16% and 8% respectively.

When accounting for the synthetic augmentation, we only increased the number of instances of the classes of interest to approximately *one eighth* the number of instances that DTLD provides for more common classes such as green solid. Further synthetic augmentation would surely yield even larger gains.

Class	Yellow Left Arrow	Red-Yellow Left Arrow
AP : Real + Synthetic	0.72	0.91
AP : Real Only	0.62	0.84

Training of Traffic Signal Classifiers is but one example of the use of highly realistic synthetic data for training ADAS and Autonomous Vehicle perception systems, and cost-effectively accelerating time to market and improving safety. Cognata provides training data for a wide variety of additional applications such as Traffic Sign, Vehicle, and Pedestrian and Cyclist Vulnerable Road User detection.

We look forward to sharing the results of other experiments in the coming months. In the meantime, visit cognata.com to learn more about our synthetic training data offerings or drop us a note to request a free evaluation dataset.

ABOUT COGNATA

Cognata is a leading global supplier of large-scale automotive simulation and synthetic training data for the Advanced Driver Assistance System (ADAS) and autonomous vehicle markets. Working with leading automotive technology companies around the world, our end to end platform accelerates time to market by delivering solutions for the entire automated driving product lifecycle, from training to testing to deployment.

cognata.com | info@cognata.com