


# Autonomous *Safety in Sight*



## By *Ansys Advantage* Staff

For a deployed autonomous vehicle (AV), there can be no surprises. Along the road or in the air, the vehicle's perception system must "make sense" of each object it "sees." For this to happen, its software models must be properly trained. Without this training, it will fail to detect or correctly classify objects it hasn't seen before.

As an example, consider the case of a person in a costume crossing the street: A human driver, although surprised, will immediately recognize the person in costume and respond accordingly. In contrast, a perception system may fail to make this critical leap in logic or, worse, fail to detect an object at all. To ensure safe operation, developers must not only train the vehicle's AI-based perception algorithm, but also ensure that it has learned what it needs to know.

Solving these and other safety-related AV issues is the mission of Pittsburgh-based Edge Case Research (Edge Case). Founded in 2014, Edge Case began as a collaboration between Carnegie Mellon University (CMU) researchers Michael Wagner (CEO) and Professor Phil Koopman, Ph.D. (CTO), who shared a commitment to building safety into autonomous systems from the ground up. The company's software products and services tackle the most complex machine learning challenges and embedded software problems. Edge Case works across the globe with original equipment manufacturers (OEMs), advanced driver assistance systems (ADAS) suppliers, Level 4+ autonomy developers, vehicle operators and insurers to help its customers go to market with products that are safe, secure and reliable.

"As we watched autonomy emerge from university research labs onto the roads, skies and hospitals, we realized we had an amazing opportunity to make autonomy safer and worthy of our trust," says Wagner.

The company's name aptly describes what it does. In the world of safety for autonomy and robotics, edge cases represent rare, potentially hazardous scenarios — and are the focus of Edge Case's product development. Switchboard, the company's initial offering, uses stress testing to automate



**Ansys SCADE Vision powered by Hologram speeds up the discovery of weaknesses in AV embedded perception software that may be tied to edge cases.**

and accelerate the finding and fixing of software defects. From Edge Case's inception, Switchboard has been an important component of the U.S. Army's efforts to improve soldier safety and advance tactical capabilities with autonomy platforms. Partnerships with other defense and autonomous technology companies, including Lockheed Martin, soon followed. Switchboard also served as an important conversation starter with Pittsburgh neighbor and soon-to-be partner, Ansys.

Edge Case's second innovation, Hologram, was conceived in 2018 as a robustness testing engine that detects weaknesses in perception systems. And, as the result of a 2019 OEM agreement, it also powers Ansys SCADE Vision, part of the Ansys product family for embedded software.

#### BEHIND THE AV SCENES IN THE NEURAL NETWORK

In an autonomous vehicle, perception is one of a number of interdependent systems — e.g., motion prediction, planning and control — that govern operation. With the exception of perception, developers had well-understood, accepted methods for ensuring the safety of these systems.

Perception developers typically use a drive-find-fix approach. Detecting defects requires looking at the output from a perception algorithm and comparing it against annotated "ground truth" object data. If, for example, the algorithm fails to detect a pedestrian, the process would be to retrain, retest and (likely) repeat.

This approach is adequate for development but insufficient for ensuring safety. First, ground truth data must be manually labeled — each object in each video frame. This is not only

**“Watching this disruptive, autonomous technology emerge from the labs around us, we recognized an amazing opportunity to make autonomy safer and worthy of the public’s trust.”**

— Dr. Phil Koopman, Edge Case Research

enormously time-consuming, but also outrageously expensive. Second, if a discrepancy is detected, the analysis wouldn’t identify the source of the weakness or why it happened. SCADE Vision not only addresses these issues, but also provides a tool for validating perception inside the larger autonomous system.

The core of a perception system is a set of sensors and a convolutional neural network (CNN). The network connects hundreds or even thousands of software-based neurons (single processing units) arranged in a series of connected layers. When processing camera sensor data from test vehicles, neurons in the input layer capture and assign numerical values to every pixel in an image. In this way, the CNN “sees” an image as an array of pixel values.

These values pass through filter-like layers that process each pixel through a series



**The tow truck in this four-image series is initially identified correctly, but additional analysis shows it could be a missed detection.**

of algebraic and matrix operations. Each layer’s decision function effectively screens for different object features — straight or curved edges, colors, textures, intensity patterns, etc. Based on the classification “decisions” made as the pixels are processed, the output layer identifies the presence of an object. It generates object lists and draws corresponding bounding boxes around pedestrians, stop signs, cars, etc.

Unlike a traditional software system, a neural network acts like a black box: It is difficult to impossible to know how it makes each decision. Edge Case created Hologram for precisely this reason. Managing the safety of perception systems is vastly different from validating rule-based planning or control systems. “There really aren’t any rules for detecting a pedestrian,” says Wagner, “and this is why handling the safety of perception systems is so different from validating rule-based planning or control systems. A neural network is built according to the training data it was fed.”





**Ansys SCADE Vision powered by Hologram looks back in time to identify edge cases that real-time monitoring alone does not detect, providing insight into the “black box” of neural networks.**

A case in point: An object detection system, which was repeatedly and correctly identifying pedestrians along a city street, failed to detect a worker at a construction site. Why? Because it had no reference in the training data for the worker's neon yellow vest. Ironically, the high-visibility vest caused the worker to disappear.

#### **ANALYZING RAW, UNLABELED DATA**

Over millions of road miles, test vehicles collect petabytes, or even exabytes, of data — of which only a fraction will be used for object detection training. This is because the data must first be labeled.

Annotation experts have to draw boxes around and label every object in every frame of video, so the CNN can learn “this is a pedestrian, this is a car.” Approximately 800 human hours are required to label just one hour of driving footage — a resource-draining, error-prone proposition.

SCADE Vision powered by Hologram doesn't require labeled data to identify fragilities in a neural network. Once a perception algorithm has been “satisfactorily” trained, SCADE Vision's automated analysis begins by running raw video footage through the neural network (referred to as the system under test, or SUT). Then it modifies the video scene, ever so slightly. The image may be blurred or sharpened, but not to the degree that a human wouldn't recognize the altered objects. Once again, SCADE Vision runs the modified frames through the SUT and compares the modified and baseline results frame by frame.

Edge Case Research's product manager, Eben Myers, says the heart of SCADE Vision's automated analysis resides in the comparison of the two sets of object detections. This is where the software engine detects edge cases and reduces petabytes of unlabeled test data to a significantly smaller subset of video frames meriting further investigation.

While minor disturbances between the detections are expected, larger disturbances (weak detections) are predictive of potential errors in the SUT (false negatives). Weak object detections signal that the software brain is straining to make a positive identification — and making a best-guess decision. False negatives, on the other hand, indicate a missed detection, an actual failure.

SCADE Vision outputs these results in two corresponding displays. A chart of the trip segment (sequential frames of video data) displays the baseline detections as gray bars, weak detections as orange and false negatives as red. And, within each frame, similarly colored (plus green for baseline detections) bounding boxes surround objects detected or missed by the SUT. Analysts can learn significantly more from SCADE Vision than from real-time analyses, as it intelligently reveals object detections in the past and in the future. For instance, a bounding box that

flickers orange in a few frames before turning (and remaining) green alerts the analyst that the SUT is confused for some reason when the object enters the scene. Without this look back, the analyst would have had no way of knowing that the green-boxed object was detected inconsistently.

SCADE Vision also finds systemic errors (versus one-offs). These present, for example, as a stop sign that shows up in several scenes, becomes a weak detection and then disappears before reappearing and repeating the pattern. This detection error could be caused by one or more triggering conditions — environmental, root-cause factors, such as glare, low contrast, noisy background (leaves on a tree), etc. SCADE Vision provides analysis tools for identifying these triggering conditions, which when combined with a weakness in the SUT, create the type of unsafe behavior the software can flag.

With these tools, analysts can add descriptive tags to the object data that characterize suspected triggering conditions and help reveal detection defect patterns. Analysts can use the output of the tagging process to perform quantitative analyses spelled out by the Safety of the Intended Functionality (SOTIF) or do additional testing on “objects of interest” in the unlabeled data. By “pointing to” these objects, SCADE Vision allows analysts to do scenario testing to see, for example, if stop signs in front of leafy trees are consistently detected. This process provides greater insight into a suspected system weakness and facilitates the retraining of the algorithm.

#### DATA LABELING BY THE NUMBERS

>24 TB data/day are generated from a 5-camera AV setup

800 human hours yield one hour of labeled data

100,000 images + one week of AI training are needed for software to learn a single traffic situation

1 M frames labeled/month are required for full-scale AV development

#### SCALING UP AV DEVELOPMENT

SCADE Vision powered by Hologram from Edge Case Research offers scalability to customers developing autonomous technology and integration with other Ansys software. It can be paired with Ansys medini analyze to discover and track identified triggering conditions. And, when used in conjunction with [Ansys VRXPERIENCE](#) and [Ansys optiSLang](#), it can automate robustness testing of the SUT over a very large number of scenario variants. SCADE Vision will also be used in conjunction with [Ansys Cloud](#) and high-performance computing (HPC) to speed the detection of edge cases at scale.

Working in partnership on perception, Edge Case and Ansys are continuing to accelerate the timeline for the safe, widespread deployment of fully autonomous vehicles.

SCADE Vision’s edge case detection capability is advancing the development of perception algorithms, while delivering a 30-times speedup in detection discovery (versus manual data analysis). What’s more, SCADE Vision can be scaled to industries beyond automotive — to mining, aerospace and defense, industrial robotics, or any application that relies on AI-based vision and perception software. ▲

