



PERCEPTION INNOVATION

Resolving Edge Cases in ADAS & Autonomous Driving

Edge Case: Abrupt Stop Detection



Introduction

Human drivers confront and handle an incredible variety of situations and scenarios—terrain, roadway types, traffic conditions, weather conditions—for which autonomous vehicle technology needs to navigate both safely, and efficiently. These are edge cases, and they occur with surprising frequency. In order to achieve advanced levels of autonomy or breakthrough ADAS features, these edge cases must be addressed. In this series, we explore common, real-world scenarios that are difficult for today's conventional perception solutions to handle reliably. We'll then describe how AEye's software definable **iDAR™ (Intelligent Detection and Ranging)** successfully perceives and responds to these challenges, improving overall safety.

Challenge: A Child Runs into the Street Chasing a Ball

A vehicle equipped with an advanced driver assistance system (ADAS) is cruising down a leafy residential street at 25 mph on a sunny day with a second vehicle following behind. Its driver is distracted by the radio. Suddenly, a small object enters the road laterally. At that moment, the vehicle's perception system must make several assessments before the vehicle path controls can react. What is the object, and is it a threat? Is it a ball or something else? More importantly, is a child in pursuit? Each of these scenarios require a unique response. It's imperative to brake or swerve for the child. However, engaging the vehicle's brakes for a lone ball is unnecessary and even dangerous.





How Current Solutions Fall Short

According to a recent study done by AAA, today's advanced driver assistance systems (ADAS) will experience great difficulty recognizing these threats or reacting appropriately. Depending on road conditions, their passive sensors may fail to detect the ball and won't register a child until it's too late. Alternatively, vehicles equipped with systems that are biased towards braking will constantly slam on the brakes for every soft target in the street, creating a nuisance or even causing accidents.

Camera. Camera performance depends on a combination of image quality, Field-of-View, and perception training. While all three are important, perception training is especially relevant here. Cameras are limited when it comes to interpreting unique environments because everything is just a light value. To understand any combination of pixels, AI is required. And AI can't invent what it hasn't seen. In order for the perception system to correctly identify a child chasing a ball, it must be trained on every possible permutation of this scenario, including balls of varying colors, materials, and sizes, as well as children of different sizes in various clothing. Moreover, the children would need to be trained in all possible variations—with some approaching the vehicle from behind a parked car, with just an arm protruding, etc. Street conditions would need to be accounted for, too, like those with and without shade, and sun glare at different angles. Perception training for every possible scenario may be possible. However, it's an incredibly costly and time-consuming process.

Radar. Radar's basic flaw is that it can only pick up a few degrees of angular resolution. When radar picks up an object, it will only provide a few detection points to the perception system to distinguish a general blob in the area. Moreover, an object's size, shape, and material will influence its detectability. Radar can't distinguish soft objects from other objects, so the signature of a rubber or leather ball would be close to nothing. While radar would detect the child, there would simply not be enough data or time for the system to detect, and then classify and react.

Camera + Radar. A system that combines radar with a camera would have difficulty assessing this situation quickly enough to respond correctly. Too many factors have the potential to negatively impact their performance. The perception system would need to be trained for the precise scenario to classify exactly what it was "seeing." And the radar would need to detect the child early enough, at a wide angle, and possibly from behind parked vehicles (strong surrounding radar reflections), predict its path, and act. In addition, radar may not have sufficient resolution to distinguish between the child and the ball.

LiDAR. Conventional LiDAR's greatest value in this scenario is that it brings automatic depth measurement for the ball and the child. It can determine within approximately a few centimeters exactly how far away each is in relation to the vehicle. However, today's LiDAR systems are unable to ensure vehicle safety because they don't gather important information—such as shape, velocity, and trajectory—fast enough. This is because conventional LiDAR systems are passive sensors that scan everything uniformly in a fixed pattern and assign every detection an equal priority. Therefore, it is unable to prioritize and track moving objects, like a child and a ball, over the background environment, like parked cars, the sky, and trees.





Successfully Resolving the Challenge with iDAR

AEye's iDAR solves this challenge successfully because it can prioritize how it gathers information and thereby understand an object's context. As soon as an object moves into the road, a single LiDAR detection will set the perception system into action. First, iDAR will cue the camera to learn about its shape and color. In addition, iDAR will define a dense **Dynamic Region of Interest (ROI)** on the ball. The LiDAR will then interrogate the object, scheduling a rapid series of shots to generate a dense pixel grid of the ROI. This dataset is rich enough to start applying perception algorithms for classification, which will inform and cue further interrogations.

Having classified the ball, the system's intelligent sensors are trained with algorithms that instruct them to anticipate something in pursuit. At that point, the LiDAR will then schedule another rapid series of shots on the path behind the ball, generating another pixel grid to search for a child. iDAR has a unique ability to intelligently survey the environment, focus on objects, identify them, and make rapid decisions based on their context.

Software Components

Computer Vision. iDAR is designed with computer vision, creating a smarter, more focused LiDAR point cloud that mimics the way humans perceive the environment. In order to effectively "see" the ball and the child, iDAR combines the camera's 2D pixels with the LiDAR's 3D voxels to create **Dynamic Vixels**. This combination helps the AI refine the LiDAR point clouds around the ball and the child, effectively eliminating all the irrelevant points and leaving only their edges.

Cueing. A single LiDAR's detection on the ball sets the first cue into motion. Immediately, the sensor flags the region where the ball appears, cueing the LiDAR to focus a Dynamic ROI on the ball. Cueing generates a dataset that is rich enough to apply perception algorithms for **classification**. If the camera lacks data (due to light conditions, etc.), the LiDAR will cue itself to increase the point density around the ROI. This enables it to gather enough data to classify an object and determine whether it's relevant.

Feedback Loops. Once the ball is detected, **a feedback loop is generated** by an algorithm that triggers the sensors to focus another ROI immediately behind the ball and to the side of the road to capture anything in pursuit, initiating faster and more accurate classification. This starts another cue. With that data, the system can classify whatever is behind the ball and determine its true velocity so that it can decide whether to apply the brakes or swerve to avoid a collision.

The Value of AEye's iDAR

LiDAR sensors embedded with AI for intelligent perception are vastly different than those that passively collect data. After detecting and classifying the ball, iDAR will immediately **foveate** in the direction where the child will most likely enter the frame. This ability to intelligently understand the context of a scene enables iDAR to detect the child quickly, calculate the child's speed of approach, and apply the brakes or swerve to avoid collision. To speed reaction times, each sensor's data is processed intelligently at the edge of the network. Only the most salient data is then sent to the domain controller for advanced analysis and path planning, ensuring optimal safety.

