Rethinking the three Rs of LiDAR

Aravind Ratnam, VP of Product Management, AEye, proposes new automotive LiDAR evaluation metrics—rate, resolution, and range—adapted to meet the capabilities of today's technology.

Aravind Ratnam leads product management at AEye, where he is responsible for product strategy, requirements, roadmap, content, and specifications. A former management consultant, Ratnam has a deep understanding of business models and broad exposure to various markets, including semiconductor, industrial and robotics, and aerospace and defense.

While once sufficient, conventional metrics used for evaluating automotive LiDAR performance must be updated to address the capabilities of modern perception systems. Today's systems are smarter and more agile—requiring a new lens through which to measure performance specifications and system effectiveness.

AEye has come up with modified metrics we believe are best suited to this task. They take the traditional "Rs" of LiDAR measurement—rate, resolution, and range—and adapt them to meet the capabilities of today's technology: Frame rate becomes intra-frame object revisit rate, resolution becomes instantaneous enhanced resolution, and detection range becomes object classification range. In this article, we will dive into all three.

Conventional metric #1: frame rate of 10 to 20 Hz

New metric: object revisit rate. (The time between two shots at the same point or set of points)

Typical LiDAR systems require multiple interrogations/detections on the same point or object to validate an object or scene. The time it takes to detect the point or object is dependent on many variables (distance, interrogation pattern and resolution, reflectivity, shape), and it traditionally takes several full frames to achieve.

What is missing from the conventional metric is a finer definition of time. AEye proposes that object revisit rate is a more refined metric for automotive LiDAR than frame rate because agile LiDAR is able to revisit an object within the same frame. Shorter object revisit times help keep processing times low, and the use of agile LiDAR accelerates revisit rate by allowing for intelligent shot scheduling within a frame while interrogating a position or object multiple times.

These interrogations can also be data dependent. For example, an object can be revisited if a (low confidence) detection occurs, and it's desirable to quickly validate (or reject) the detection with a secondary measurement, as seen in Figure 1.

An accelerated object revisit rate, therefore, allows for faster data acquisition because it can identify and automatically revisit an object, painting a more complete picture of it within the context of the scene, such as identifying a pedestrian next to oncoming headlights or a semi-trailer laterally crossing the path of the vehicle.

Conventional metric #2: fixed (angular) resolution over a fixed field-of-view

New metric: instantaneous (angular) resolution. (The degree to which a LiDAR sensor can apply additional resolution to key areas within a frame)

The conventional metric of fixed resolution assumes that the field-of-view will be scanned with a constant pattern (which is true for unintelligent sensors that have limited or no ability to adapt their collection capabilities), and that salient information resident within the scene is uniform in space and time (which is not true).

Instantaneous resolution allows for the creation of multiple regions-of-interest within a scene, allowing the system to focus and gather more comprehensive data about specific objects. For example, if the system encounters a jaywalking pedestrian directly in the path of the vehicle, it can focus more of its attention on the jaywalker and less on parked vehicles along the side of the road.

AEye modeled its iDAR after the world's most intelligent and efficient perception system, the human visual cortex, which filters out irrelevant information, such as houses or trees, while simultaneously focusing on a particular point of interest, like a busy intersection ahead. This is called foveation, where the target of a gaze is allotted a higher concentration of retinal cones, allowing it to be seen more vividly while other, less important objects are pushed to the periphery.

Figure 2 shows the limitations of conventional LiDAR, whereas the company's solution enables dynamic perception, allowing the system to focus on and gather more comprehensive data about a particular ROI (region of interest) at fast speed.

Conventional metric #3: object detection range

New metric: object classification range. (Range at which you have sufficient data to classify an object)

Assuming that the purpose of the conventional metric is to accurately score how LiDAR systems contribute to autonomous vehicle safety, automotive LiDAR evaluators should also consider how long it takes these systems to identify hazards. Thus, object classification range becomes a far more meaningful metric because what matters most in autonomous driving is how quickly an object can be identified and classified, a threat-level decision made, and an appropriate response calculated.

Currently, classification typically takes place in the perception stack. It's at this point that objects are labeled and, eventually, more clearly identified. AEye argues that determining a system's object classification range reduces the unknowns—such as latency associated with noise suppression—early in the perception stack, pinpointing the salient information about whether a LiDAR system is capable of operating at optimal safety.

As a relatively new field, the definition of how much data is necessary for classification in automotive LiDAR has not yet been defined. Thus, we propose that adopting perception standards used by video classification provides a valuable provisional definition. According to video standards, enabling classification begins with a 3 x 3 pixel grid of an object. Under this definition an automotive LiDAR system might be assessed by how fast it's able to generate a high-quality, high-resolution 3 x 3 point cloud that enables the perception stack to comprehend objects and people in a scene.

In Figure 3, we compare this scenario (Scan 1) to one in which a company sensor (Scan 2), detects the object ahead.

The fundamental value of being able to classify objects at range is greatest in instances where the identity of the object defines a specific and immediate response from the vehicle, such as when encountering a school bus full of children. The faster that object is classified as a school bus, the faster the autonomous vehicle can initiate an appropriate response.

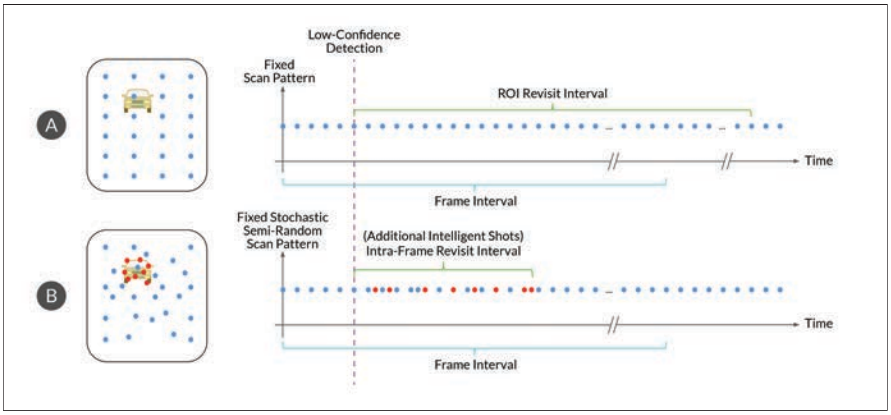

Figure 1. Advanced agile LiDAR sensors enable intelligent scan patterns such as the "foveation in time" intra-frame revisit interval and random scan pattern of the AEye solution (B) compared to revisit interval on a typical fixed pattern LiDAR (A).

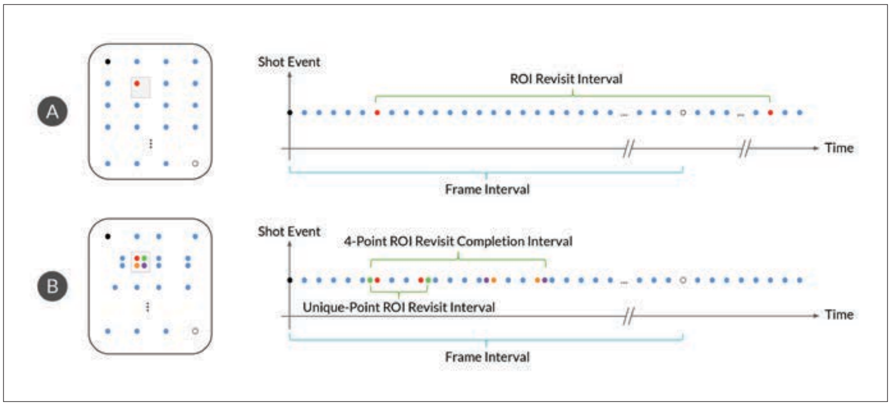

Figure 2. Region of interest (ROI) revisit rate and foveation of the AEye solution (B) compared to conventional scan patterns (A).

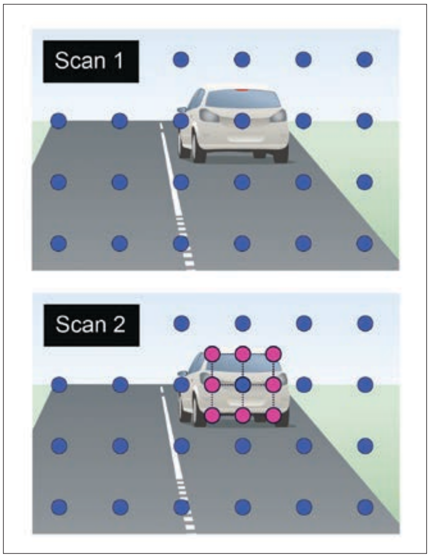

Figure 3. Packing a dense 3 x 3 grid around a detect allows the collection of more useful data and greatly speeds up classification. Scan 1 above has a single detect on a vehicle. Rather than wait for the next frame to resample this vehicle (as is the traditional mode in LiDAR) we instead quickly form a dedicated dense ROI, as indicated in Scan 2 below. This is done almost immediately after the initial single detect and before completing the next scan.