Rethinking the Four "Rs" of LiDAR: Rate, Resolution, Returns and Range

Extending Conventional LiDAR Metrics to Better Evaluate Advanced Sensor Systems

By Blair LaCorte, Luis Dussan, Allan Steinhardt, and Barry Behnken

RRRRWP_21_1026

1 Park Place, Suite 200, Dublin, CA 94568 • 925.400.4366

Executive Summary

As the autonomous vehicle market matures, sensor and perception engineers have become increasingly sophisticated in how they evaluate system efficiency, reliability, and performance. Many industry leaders have recognized that conventional metrics for LiDAR data collection (such as frame rate, full frame resolution, points per second, and detection range) no longer adequately measure the effectiveness of sensors to solve real-world use cases that underlie autonomous driving.

First generation LiDAR sensors passively search a scene and detect objects using background patterns that are fixed in both time (no ability to enhance with a faster revisit) and in space (no ability to apply extra resolution to high interest areas like the road surface or pedestrians). A new class of solid-state, high-performance, active LiDAR sensors enable intelligent information capture that expands their capabilities—moving from "passive search" or detection of objects, to "active search," and in many cases, to the actual acquisition of classification attributes of objects in real time.

Because early generation LiDARs use passive fixed raster scans, the industry adopted very simplistic performance metrics that don't capture all the nuances of the sensor requirements needed to enable AVs. In response, AEye is proposing the consideration of four new corresponding metrics for extending LiDAR evaluation. Specifically: extending the metric of *frame rate* to include ***object revisit rate;*** extending the metric of *resolution* to capture ***instantaneous resolution;*** extending *points per second* to signify the overall more useful ***quality returns per second;*** and extending *detection range* to reflect the more critically important ***object classification range.***

We are proposing that these new metrics be used in conjunction with existing measurements of basic camera, radar, and passive LiDAR performance. These extended metrics measure a sensor's ability to intelligently enhance perception and create a more complete evaluation of a sensor system's efficacy in improving the safety and performance of autonomous vehicles in real-world scenarios.

Introduction

Our industry has leveraged proven frameworks from advanced robotic vision research and applied them to LiDAR-specific product architectures. One framework, ***"Search, Acquire [or classify], and Act,"*** has proven to be both versatile and instructive relative to object identification.

- **Search** is the ability to detect any and all objects without the risk of missing anything.
- **Acquire** is defined as the ability to take a search detection and enhance the understanding of an object's attributes to accelerate classification and determine possible intent (this could be done by classifying object type or by calculating velocity).
- **Act** defines an appropriate sensor response as trained, or as recommended, by the vehicle's perception system or domain controller. Responses can largely fall into four categories:
  - Continue scan for new objects with no enhanced information required;
  - Continue scan and interrogate the object further, gathering more information on an acquired object's attributes to enable classification;
  - Continue scan and track an object classified as non-threatening;
  - Continue scan and instruct the control system to take evasive action.

Within this framework, performance specifications and system effectiveness need to be assessed with an "eye" firmly on the ultimate objective: completely safe operation of the vehicle. However, as most LiDAR systems today are passive, they are only capable of basic search. Therefore, conventional metrics used for evaluating these systems' performance relate to basic object detection capabilities—frame rate, resolution, points per second, and detection range. If safety is the ultimate goal, then ***search*** needs to be more intelligent, and ***acquisition*** (and classification) done more quickly and accurately so that the sensor or the vehicle can determine how to ***act*** immediately.

Rethinking the Metrics

Makers of automotive LiDAR systems are frequently asked about their frame rate, and whether or not their technology has the ability to detect objects with 10% reflectivity at some range (often 230 meters). We believe these benchmarks are required, but insufficient as they don't capture critical details, such as the size of the target, the speed at which it needs to be detected and recognized, or the cost of collecting that information.

We believe it would be productive for the industry to adopt a more holistic approach when it comes to assessing LiDAR systems for automotive use. We argue that we must look at metrics as they relate to a perception system in general, rather than as an individual point sensor, and ask ourselves: ***"What information would enable a perception system to make better, faster decisions?"*** In this white paper, we outline the four conventional LiDAR metrics with recommendations on how to extend them.

Conventional Metric #1: Frame Rate of 10Hz – 20Hz

Extended Metric: Object Revisit Rate

The time between two shots at the same point or set of points

Defining single point detection range alone is insufficient because a single interrogation point (shot) rarely delivers sufficient confidence—it is only suggestive. Therefore, passive LiDAR systems need either multiple interrogations/detects at the same location or multiple interrogations/detects on the same object to validate an object or scene. In passive LiDAR systems, the time it takes to detect an object is dependent on many variables, such as distance, interrogation pattern, resolution, reflectivity, the shape of the object, and the scan rate.

A key factor missing from the conventional metric is a finer definition of time. Thus, we propose that **object revisit rate** become a new, more refined metric for automotive LiDAR because a high-performance, active LiDAR, such as AEye's iDAR™, has the ability to revisit an object within the same frame. The time between the first and second measurement of an object is critical, as shorter object revisit times keep processing times low for advanced algorithms that correlate multiple moving objects in a scene. The best algorithms used to associate/correlate multiple moving objects can be confused when time elapsed between samples is high. This lengthy combined processing time, or latency, is a primary issue for the industry.

The active iDAR platform accelerates revisit rate by allowing for intelligent shot scheduling within a frame. Not only can iDAR interrogate a position or object multiple times within a conventional frame, it can maintain a background search pattern while simultaneously overlaying additional intelligent shots. For example, an iDAR sensor can schedule two repeated shots on an object of interest in quick succession (30μsec). These multiple interrogations can be contextually integrated with the needs of the user (either human or computer) to increase confidence, reduce latency, or extend ranging performance.

These additional interrogations can also be data dependent. For example, an object can be revisited if a low confidence detection occurs, and it is desirable to quickly validate or reject it, enabled with secondary data and measurement, as seen in *Figure 1*. A typical frame rate for conventional passive sensors is 10Hz. For conventional passive sensors, this is the object revisit rate. With AEye's active iDAR technology, the object revisit rate is now different from the frame rate, and it can be as low as tens of microseconds between revisits to key points/objects— easily 100x to 1000x faster than conventional passive sensors.

What this means is that a perception engineering team using dynamic object revisit capabilities can create a perception system that is at least an order of magnitude faster than what can be delivered by conventional passive LiDAR without disrupting the background scan patterns. We believe this capability is invaluable for delivering level 4/5 autonomy as the vehicle will need to handle complex edge cases, such as identifying a pedestrian in front of oncoming headlights or a flatbed semi-trailer laterally crossing the path of the vehicle.

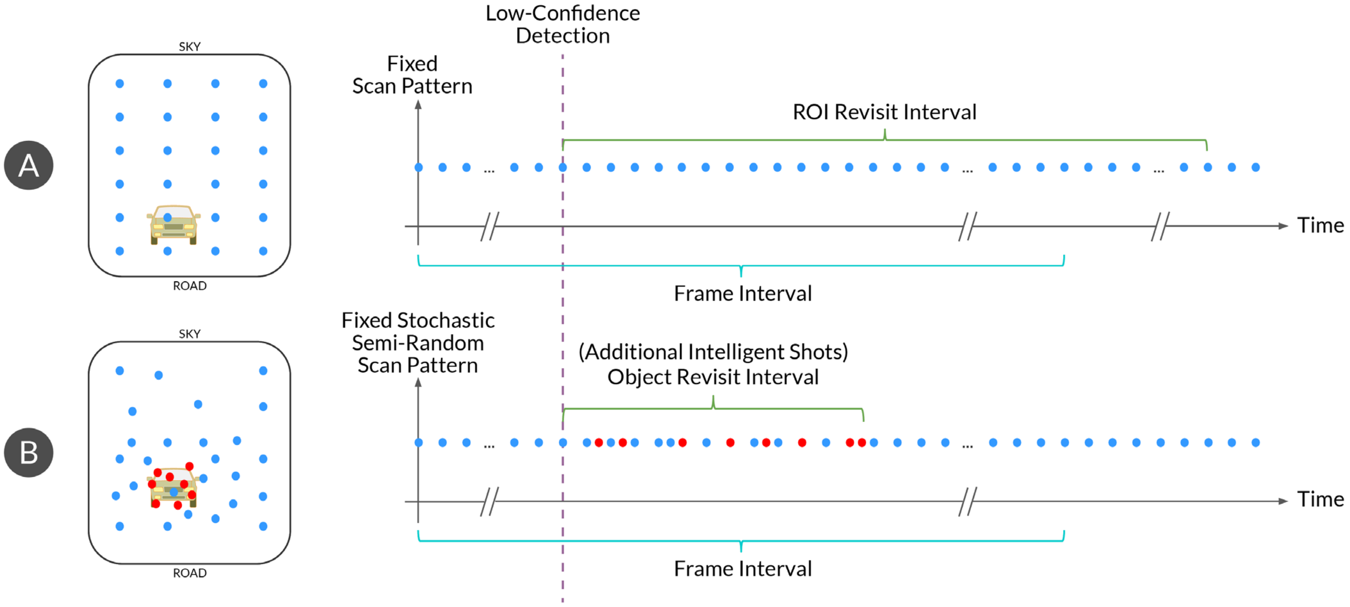

Figure 1. Advanced active LiDAR sensors utilize intelligent scan patterns that enable an Object Revisit Interval, such as the random scan pattern of AEye's iDAR (B). This is compared to the Revisit Interval on a passive, fixed pattern LiDAR (A). For example, in this instance, iDAR is able to get eight detects on a vehicle, while passive, fixed pattern LiDAR can only achieve one.

Within the "Search, Acquire, and Act" framework, an accelerated **object revisit rate,** therefore, allows for faster acquisition because it can identify and automatically revisit an object, painting a more complete picture of it within the context of the scene. Ultimately, this allows for collection of object classification attributes in the sensor, as well as efficient and effective interrogation and tracking of a potential threat.

Real-World Applications

Use Case: Head-On Detection

When you're driving, the world can change dramatically in a tenth of a second. In fact, two cars traveling towards each other at 100 kph are 5.5 meters closer after 0.1 seconds. By having an accelerated revisit rate, we increase the likelihood of hitting the same target with a subsequent shot due to the decreased likelihood that the target has moved significantly in the time between shots. This helps the user solve the "Correspondence Problem," determining which parts of one "snapshot" of a dynamic scene correspond to which parts of another snapshot of the same scene. It does this while simultaneously enabling the user to quickly build statistical measures of confidence and generate aggregate information that downstream processors might require, such as object velocity and acceleration. The ability to selectively increase revisit rate on objects of interest while lowering the revisit rate in sparse areas, like the sky, can significantly aid higher level inferencing algorithms, allowing perception and path planning systems to more quickly determine optimum autonomous decision making.

Use Case: Lateral Detection

A vehicle entering a scene laterally is the most difficult to track. Even Doppler Radar has a difficult time with this scenario. However, selectively allocating shots to extract velocity and acceleration when detections have occurred as part of the acquisition chain vastly reduces the required number of shots per frame. Adding a second detection, via iDAR, to build a velocity estimate on each object detection increases the overall number of shots by only 1%. Whereas, obtaining velocity everywhere with a fixed scan system doubles the required number of shots. This speed and shot saliency makes autonomous driving much safer because it eliminates ambiguity and allows for more efficient use of processing resources.

The AEye Advantage

Whereas other LiDAR systems are limited by the physics of fixed laser pulse energy, fixed dwell time, and fixed scan patterns, iDAR is a software-configurable system that allows perception and motion planning modules to dynamically customize their data collection strategies to best suit their information processing needs at design time and/or run time.

iDAR's unique bore-sighted design eliminates parallax between the camera and the LiDAR, bringing it extremely close to solving the "Correspondence Problem." The achievable object revisit rate of AEye's iDAR system for points of interest (not merely the exact point just visited) is microseconds to a few milliseconds—which can be up to 3000x faster, compared to conventional LiDAR systems that typically require hundreds of milliseconds between revisits. This gives the unprecedented ability to calculate valuable attributes such as object velocity (both lateral and radial) faster than any other system, allowing the vehicle to act more readily to immediate threats and track them through time and space more accurately.

This ability to define the new metric, **object revisit rate,** which is decoupled from the traditional "frame rate," is important also for the next metric we introduce. This second metric helps to distinguish "search" algorithms from "acquisition" algorithms. Separation of these two types of algorithms provides insight into the heart of iDAR, which is the principle of information quality (as opposed to data quantity): "more information, less data."

Conventional Metric #2: Fixed Resolution Over a Fixed Field-of-View

Extended Metric: Instantaneous Resolution

The degree to which a LiDAR sensor can apply additional resolution to key areas within a frame

Resolution as a conventional metric assumes that the Field-of-View will be scanned with a constant pattern and with uniform power. This makes perfect sense for less intelligent, passive sensors that have a limited ability to adapt their collection capabilities. Additionally, the conventional metric assumes that salient information within the scene is uniform in space and time, which we know is not true. This is especially apparent for a moving vehicle. However, because of these assumptions, conventional LiDAR systems indiscriminately collect gigabytes of data from a vehicle's surroundings, sending those inputs to the CPU for decimation and interpretation.

An estimated 75% to 95% of this data is found to be useless or redundant and thrown out. In addition, these systems apply the same level of power everywhere, such that the sky is scanned at the same power as an object directly in the path of the vehicle. It's an incredibly inefficient process.

As humans, we don't "take in" everything around us equally. Rather, our visual cortex filters out irrelevant information, such as an airplane flying overhead, while simultaneously (not serially) focusing our eyes on a particular point of interest. Focusing on a point of interest allows other, less important objects to be pushed to the periphery. This is called foveation, where the target of our gaze is allotted a higher concentration of retinal cones, thus allowing it to be seen more vividly.

iDAR uses biomimicry (see the AEye white paper, ***The Future of Autonomous Vehicles: Think Like a Robot, Perceive Like a Human***) to apply and expand upon the capabilities of the human visual cortex for artificial perception. Whereas humans typically only foveate on one area, iDAR can foveate on multiple areas simultaneously (and in multiple ways), while also maintaining a background scan to ensure it never misses new objects. We describe this feature as a Region of Interest (ROI). Furthermore, since humans rely entirely on light from the sun, moon, or artificial lighting, human foveation is "receive only," i.e., passive. iDAR, by contrast, foveates on both transmit (regions that the laser light chooses to "paint") and receive (where/when the processing chooses to focus).

An example of this follows.

Figure 2 shows two systems, System A and System B. Both systems have a similar number of shot points on the same scene (left). System A represents a uniform scan pattern, typical of conventional, passive LiDAR sensors. These fixed scan patterns produce a fixed frame rate with no concept of an ROI. System B shows an adjusted, active scan pattern. The shots in System B are gathered more densely within and around the ROI (the small box) within the square. In addition, the background scan continues to search to ensure no new objects are missed, while focusing additional resolution on a fixed area to aid in acquisition. In essence, it is using intelligence to optimize the use of power and shots.

Looking at the graphs (right) associated with Systems A and B, we see that the active scan pattern of System B can revisit an ROI within a much shorter interval than the fixed scan pattern of System A. System B not only can complete one ROI revisit interval, but multiple ROIs within a single frame. Whereas, System A cannot revisit. iDAR does what conventional, passive LiDAR cannot: it enables dynamic perception, allowing the system to focus on, and gather more comprehensive data about, a particular Region of Interest at unprecedented speed.

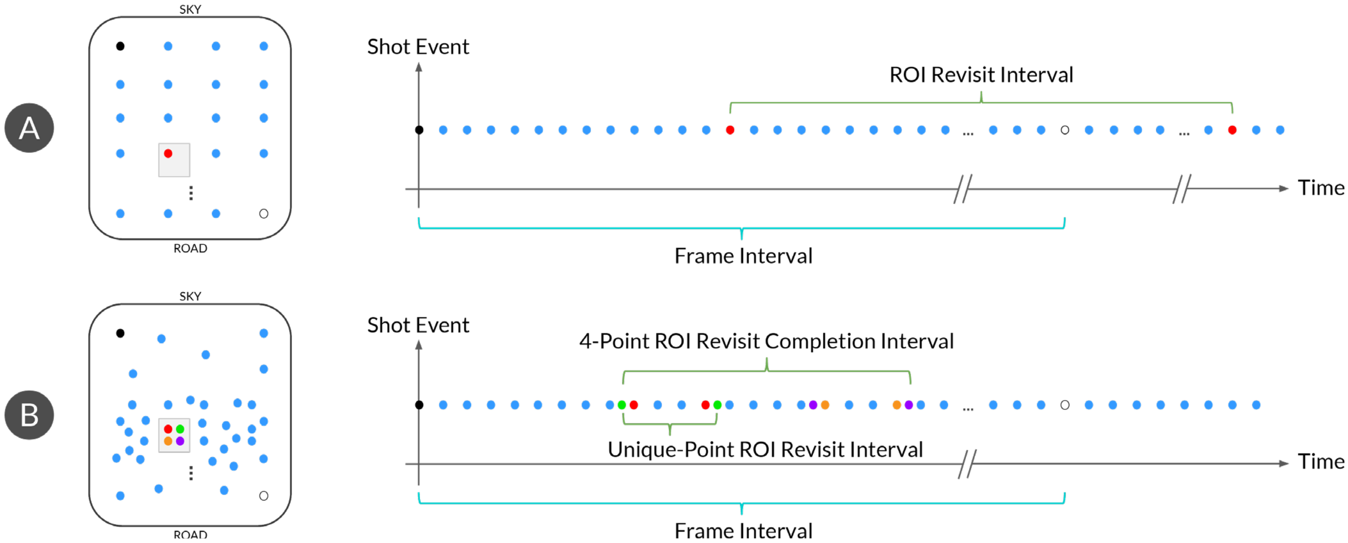

Figure 2. Region of Interest (ROI) and foveation of iDAR (B) compared to conventional scan patterns (A).

Within the "Search, Acquire, and Act" framework, instantaneous resolution allows the iDAR system to search an entire scene and acquire multiple targets, capturing additional information about them. iDAR also allows for the creation of multiple simultaneous ROIs within a scene, allowing the system to focus and gather more comprehensive data about specific objects, enabling it to interrogate them more completely and track them more effectively.

Real-World Application

Use Case: Object Interrogation

When objects of interest have been identified, iDAR can “foveate” its scanning to gather more useful information about them and acquire additional classification attributes. For example, let’s say the system encounters a jaywalking pedestrian directly in the path of the vehicle. Because iDAR enables a dynamic change in both temporal and spatial sampling density within a Region of Interest, what we call instantaneous resolution, the system can focus more of its attention on this jaywalker, and less on irrelevant information, such as parked vehicles along the side of the road. Regions of Interest allow iDAR to quickly, efficiently, and accurately identify critical information about the jaywalker, such as speed and direction. The iDAR system provides the most useful, actionable data to the domain controller to help determine the most timely course of action.

We see **instantaneous resolution** being utilized in three primary ways to address different use cases:

1. Fixed Region of Interest (ROI): Today, passive systems can only allocate more scan lines at the horizon—a very simple foveation technique limited by their fixed resolution. With second generation intelligent systems, like iDAR, that enable **instantaneous resolution,** an OEM or Tier 1 will be able to utilize advanced simulation programs to test hundreds (or even thousands) of shot patterns—varying speed, power, and other constraints—to identify an optimal pattern that integrates a fixed ROI with higher instantaneous resolution to achieve their desired results.

For example, a fixed ROI could be used to optimize the shot pattern of a unit behind a windshield with varying rakes. Additionally, a fixed ROI could be used in urban environments, where threats are more likely to come from the side of the road—such as car doors opening, pedestrians, and cross traffic—or in the immediate path of the vehicle. An ROI is defined by applying additional resolution to a fixed region that covers both sides of the road and the road surface immediately in front of the vehicle (see Figure 3B). This instantly provides superior resolution (both vertical and horizontal) in the area of greatest concern. Once a pattern is approved, it can be fixed for functional safety.

2. Triggered ROI: A Triggered ROI requires a software-configurable system that can be programmed to accept a trigger. The perception software team may determine that when certain conditions are met, an ROI is generated within the existing scan pattern. For example, a mapping or navigation system might signal that you are approaching an intersection, which generates an appropriately targeted ROI on key areas of the scene with greater detail (*see Figure 3C*).

3. Dynamic ROI: A Dynamic ROI requires the highest level of intelligence and utilizes the same techniques and methodology deployed by Automatic Targeting Systems (ATS) in fighter jets to continuously interrogate objects of high interest over time. As these objects move closer or further away, the size and density of the ROI varies. For example, pedestrians, cyclists, vehicles, or other objects moving in the scene can be detected and a Dynamic ROI automatically applied to track their movements (*see Figure 3D*).

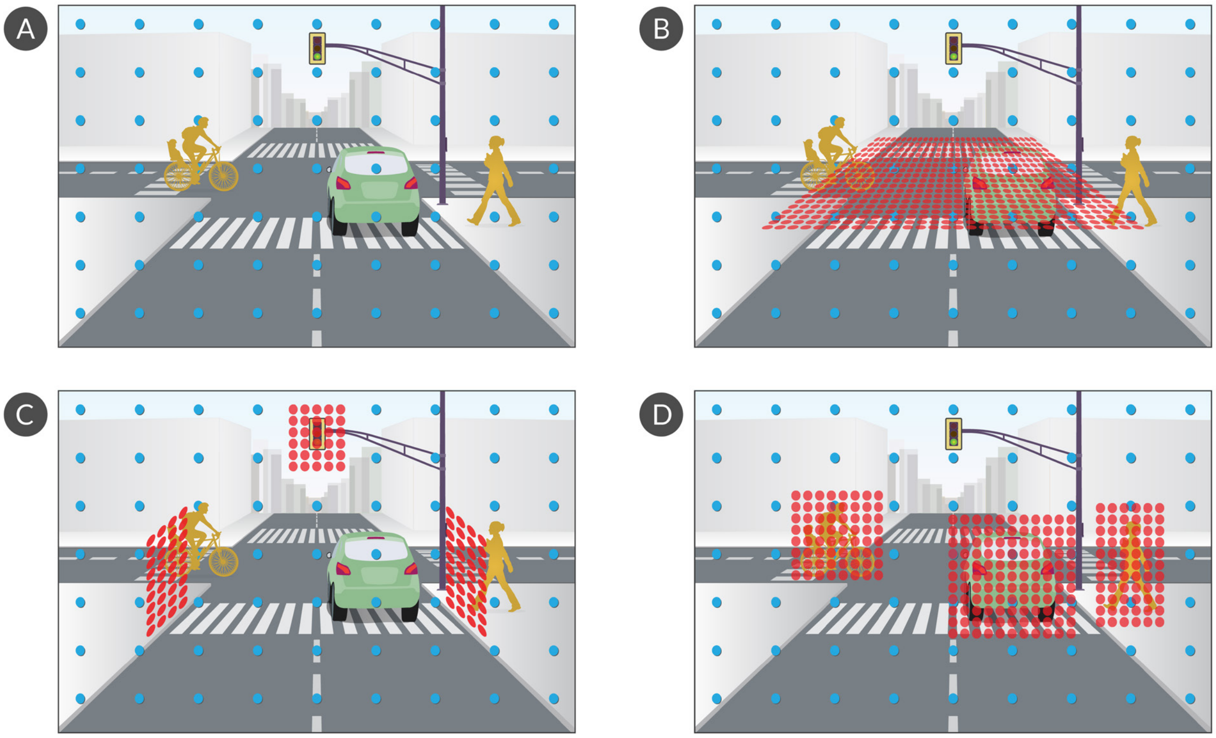

Figure 3. Figure 3A shows a scene as a vehicle approaches an intersection. Figure 3B shows a Fixed Region of Interest (ROI) covering the sides of the road and the area immediately in front of the vehicle. Figure 3C shows a Triggered ROI where the navigation system triggers specific ROIs as the vehicle approaches the intersection. Figure 3D shows a Dynamic ROI where several objects of interest are detected and tracked as they move through the scene.

The AEye Advantage

A major advantage of iDAR is that it is active in nature, meaning it can adjust its scan patterns in real time, and therefore, can take advantage of concepts like time multiplexing. This means it can simultaneously trade off temporal sampling resolution, spatial sampling resolution, and even range, at multiple points in the “frame.” This allows the system to dynamically change the scan density over the entire Field-of-View, enabling the robust collection of useful, actionable information.

In a conventional LiDAR system, there is (i) a fixed Field-of-View, (ii) a fixed uniform or patterned sampling density, and (iii) a fixed laser shot schedule. AEye’s technology allows for these three parameters to vary almost independently. This leads to an endless stream of potential innovations and will be the topic of a later paper.

Instantaneous resolution conveys that resolution, as a metric, is not something dictated by physical constraints alone, such as beam divergence, or number of points per second (the next metric). Rather, it starts with a faster, more efficient active LiDAR and then intelligently optimizes resources. The ability to instantaneously increase resolution is a critical enabler in the fourth metric we introduce.

Conventional Metric #3: Points per Second

Extended Metric: Quality Returns per Second

A high confidence, confirmatory return from an object on a single frame basis

Traditionally, the industry has favored achieving the highest number of points per second. In theory, a higher number of laser shots would mean that the sensor system would receive a higher number of returns. However, a high number of shots does not guarantee a high number of returns, nor does it necessarily mean that the data being returned is useful in any way to help safely and efficiently guide an autonomous vehicle. As mentioned earlier, passive conventional LiDAR systems simply gather data about the environment indiscriminately and without discretion, sending those inputs to the CPU, wherein, 75% to 95% of this data is thrown out. This creates a huge strain on interrogation times, bandwidth, and processing. Therefore, a conventional system that is purporting to deliver a high *quantity of shots* (i.e., high rate of shots per second) will suffer a latency penalty because it cannot separate the valuable information from the invaluable (or redundant) in a timely manner.

As safety is the ultimate goal of these systems, then having full scene coverage without missing anything, while simultaneously increasing probability of detection (i.e., knowing something is there) and reducing false positives, is a fundamental requirement. AEye proposes replacing the metric of points per second with the more meaningful quality returns per second. Measuring **quality returns per second** is significantly more beneficial to the development of automotive LiDAR systems because it quantifies the crux of the information actually needed to enable accurate and efficient perception. While points per second gives little to no indication of the value of the information received, **quality returns per second** does.

Because there is no agreed upon standard in the industry for measuring returns per second, AEye defines quality returns per second as: high probability of detection, low false positive rate, often non-isolated, returns from an object on a single frame basis (i.e., “quality” returns). And, in valuing the efficiency of LiDAR sensor systems (and thus, the safety of the autonomous vehicle and its passengers) above all else, we urge the rest of the LiDAR community to do the same.

A high probability of detection, low false positive rate return will deliver actionable data to the vehicle’s perception system. AEye is able to define quality returns per second in this way because of our bistatic architecture (the subject of a future paper). As mentioned earlier, iDAR can foveate and isolate on both transmit (regions that the laser light chooses to “paint”) and receive (where/when the processing chooses to focus). Our patented bistatic architecture keeps the transmit and receive channels separate, allowing optimization of both paths. As each laser pulse is transmitted, the receiver is told where and when to look for its return. No other LiDAR sensor system on the market can do this. Our system is agile and active enough to enable it to select where to scan, making the returns that much more efficient in capturing the most salient, actionable data.

Within the "Search, Acquire, and Act" framework, our bistatic architecture allows the iDAR system to *search* the entire scene without missing anything, but focus on what matters most in a vehicle's surroundings, actively favoring the swift *acquisition* and *tracking* of real, actionable data for smarter, more accurate decision making *(action)* and safer vehicle autonomy.

Real-World Application

Use Case: Brick in the Road on the Highway

Being able to swiftly and accurately acquire small objects on the road can be the key to preventing fatal high-speed accidents. Imagine a vehicle with highway autopilot driving at fast speeds on the highway. The vehicle detects an object close to the road. Is it a false positive? It is tumbleweed or is it a brick? Being able to acquire the object and the overridability with absolute certainty (high probability of detection, low false positive rate, and multiple returns) as quickly as possible is critical in this scenario. By favoring quality returns per second (as opposed to straining interrogation times and bandwidth on the collection of irrelevant data to receive the highest number of returns possible), low probability of detection cues the system to instantaneously interrogate further (see *Figure 4*).

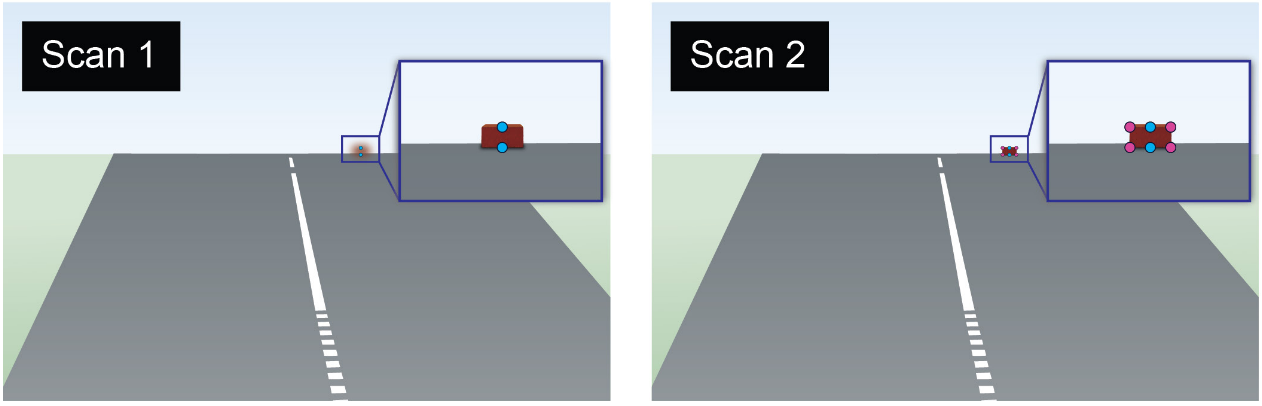

Figure 4. What is that object on the road ahead? AEye's active, intelligent iDAR is able to acquire the object as quickly as possible. By favoring quality returns per second, low probability of detection cues the system to instantaneously interrogate further.

Use Case: Child at Play

The ability to acquire and track pedestrians in complex urban or suburban environments can be a matter of life and death. Imagine an autonomous vehicle driving through a residential neighborhood when suddenly, the vehicle detects an object approaching the street. It appears to be a child chasing a ball. Being able to focus and acquire the detected object as quickly as possible is critical in this scenario. Using quality returns per second ensures that all objects are appropriately revisited, acquired, and tracked, providing salient information as quickly as possible.

The AEye Advantage

AEye's bi-static architecture enables the optimization of information in both directions, allowing for our returns to be highly curated and more informative. The AEye system is able to achieve 4 confirmatory returns at different positions (with confidence that they are different objects) from 1 shot. (Although some systems have the ability to provide multiple returns, we aren't aware of any other system meeting or exceeding this ability.) Ultimately, this ability enables the AEye system to 1) have more returns per second 2) gain more information about the object to inform perception and 3) attain more information for better and more actionable scene coverage.

To put safety and efficiency of automotive LiDAR systems first, evaluators must favor measuring high confidence returns, (i.e., **quality returns per second**) over points per second, which gives little to no indication of the value of the information received. As already stated, there is a difference between quantity and quality returns. Any LiDAR system can send out a high number of points per second, but only AEye enables the collection of the most **high-quality returns per second,** cutting down latency and enabling the safest, most accurate perception for vehicle autonomy. Our bi-static architecture enables deterministic artificial intelligence to be introduced into the sensing process at the point of acquisition. Ultimately, this establishes the iDAR platform as active—allowing it to focus on what matters most in a vehicle's surroundings (without missing anything) while still providing coverage for the entire scene. The result mimics how the human visual cortex conceptually focuses on and evaluates the environment around the vehicle, driving conditions, and road hazards, enabling smarter, more accurate decision making—radically improving the probability of detection (i.e., knowing something is there) and the accuracy of classification. This leads us to our fourth and final extended metric.

Conventional Metric #4: Object Detection Range

Extended Metric: Object Classification Range

Range at which you have sufficient data to classify an object

When it comes to measuring how well automotive LiDAR systems perceive the space around them, manufacturers commonly agree that it is valuable to determine their detection range. To optimize safety, the on-board computer system should detect obstacles as far ahead as possible. The speed with which they can do so theoretically determines whether control systems can plan and perform timely, evasive maneuvers. AEye believes that detection range is a required, but insufficient metric. Ultimately, it is the perception system's ability to classify an object and pass along accurate and timely information to the control system that enables the control system to decide on a basic course of action.

What matters most is not only how quickly an object can be detected, but how quickly it can be identified and classified, a threat-level decision made, and an appropriate response calculated. A single point detection is indistinguishable from noise. Therefore, we use a common industry definition for detection which involves persistence in adjacent shots per frame and/or across frames. We require five detects on an object per frame (five points at the same range) and/or from frame-to-frame (one single related point in five consecutive frames) to declare that a detection is a valid object. At 20Hz, it takes .25 seconds to define a simple detect.

Currently, classification typically takes place in the perception stack. It's at this point that objects are labeled and, eventually, more clearly identified. This data is used to predict behavior patterns or trajectories. The more the sensor can provide classification attributes, the faster the perception system can confirm and classify.

AEye argues that a better measurement for assessing this critical automotive LiDAR capability is its ability to impact **object classification range.** This metric reduces the unknowns—such as latency associated with noise suppression (e.g., N of M detections)—early in the perception stack, pinpointing the salient information.

As a relatively new field, the definition of how much data is necessary for classification in automotive LiDAR has not yet been specified. Thus, we propose that perception standards used by video classification offer a valuable proxy definition. According to video standards, enabling classification begins with a 3x3 pixel grid of an object. Under this definition, an automotive LiDAR system might be assessed by how fast it is able to generate a high-quality, high-resolution 3x3 point cloud that enables the perception stack to comprehend objects and people in a scene.

Generating a 3x3 point cloud is a struggle for conventional LiDAR systems. While many systems tout an ability to manifest point clouds comprising half a million or more points in one second, there is a lack of uniformity in these images. These fixed sampling patterns can be difficult for classification routines because the domain controller has to grapple with half a million points per second that are, in many cases, out of balance with the resolution required for the critical sampling of the object in question. Such a broad sample of points means it needs to perform additional interpretation, straining CPU resources.

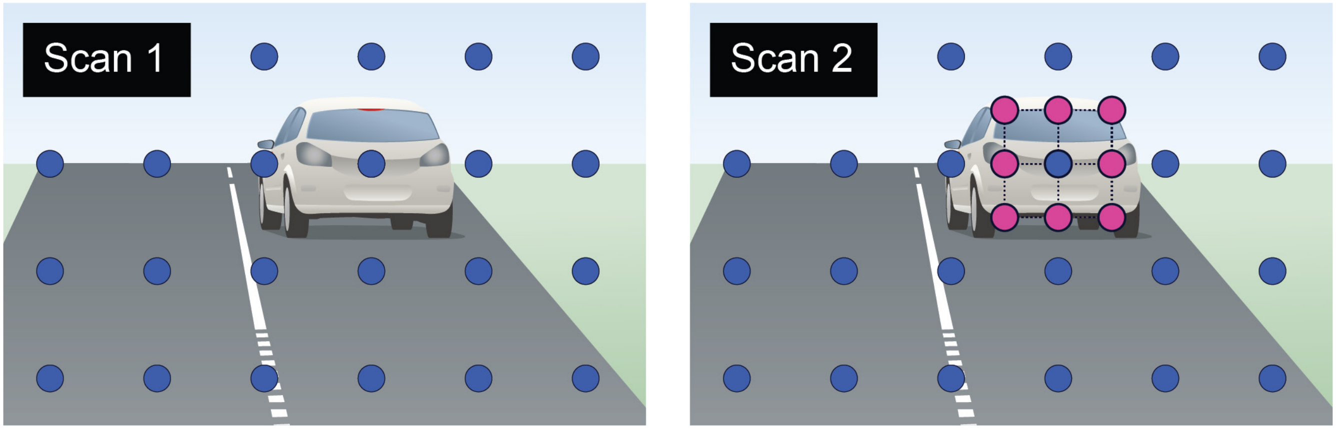

Figure 5. Packing a dense 3x3 grid around a detect allows the collection of more useful data and greatly speeds up classification. In “Scan 1” on the left, we have a single detect on a vehicle. Rather than wait for the next frame to resample this vehicle (as is typical in conventional passive LiDAR) we instead quickly form a Dynamic ROI (as seen in “Scan 2”). This is done immediately after the initial single detect, and before completing the next scan.

Returning to the “Search, Acquire, and Act” framework, once we have *acquired* a target and determined that it is valid and a possible threat, we can allocate more shots for classification and take *action* if need be. Alternatively, if we determine that the target is not an immediate threat, we can more fully interrogate the object for additional classification data or simply *track* it with a few shots per scan.

Real-World Applications

Use Case: Unprotected Left-Hand Turn

Different objects demand different responses. This is especially true in challenging driving scenarios such as an unprotected left-hand turn—especially when traversing across high-speed, oncoming traffic. Imagine an autonomous vehicle on a four-lane road with a speed limit of 100 kph needing to make an unprotected left-hand turn across two lanes of traffic. In the oncoming traffic, one lane has a motorcycle and the other has a car. In this situation, object classification range is critical, as classifying one of the objects as a motorcycle at sufficient range would indicate that the autonomous vehicle should behave more cautiously in proceeding, as motorcycles are capable of traveling at higher speeds and can take more unpredictable paths.

Use Case: School Bus (Unique Object)

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. An excellent example of this is encountering a school bus full of children. The faster that object is classified specifically as a school bus, the faster the autonomous vehicle can initiate an appropriate protocol—slowing the vehicle and deploying other tools, such as **instantaneous resolution (Triggered ROIs),** in areas around the school bus to immediately capture any movement of children toward the path of the vehicle. This capability enables similarly specific responses for police cars, ambulances, fire trucks, or any vehicle that would require the autonomous vehicle to alter how it interrogates the scene and/or change its speed or path.

The AEye Advantage

LiDAR sensors embedded with deterministic AI for perception are very different from those that passively collect data. AEye's active system can *acquire* targets and enable classification in far less time than conventional passive LiDAR systems would require to merely register a detection. With the ability to modulate revisit rate up to 3000x faster in a frame, iDAR no longer focuses on detection alone: it is now more important to gauge speed of *acquisition* (i.e., classification range). This brings to light the difference between detection range and **object classification range.**

Assuming that metrics like detection range are used for accurately scoring how LiDAR systems contribute to autonomous vehicle safety, then evaluators should also consider how long it takes these systems to *identify hazards.* Thus, **object classification range** is a far more meaningful metric.

Conclusion

In this white paper, we have discussed why reducing the time between object detections within the same frame is critical. As capturing multiple detects of the same point/object is required to fully comprehend an object or scene, measuring **object revisit rate** is a more critical metric for automotive LiDAR than frame rate.

In addition, we have argued that quantifying resolution is insufficient. It is more important to quantify **instantaneous resolution** because intelligent, agile resolution is more efficient and provides greater safety through faster response times, especially when pairing ROIs with convolutional neural networks (the subject of a future paper).

We have also demonstrated why the conventional metric of points per second is inadequate to the safety and efficiency of autonomous vehicles, and that measuring **quality returns per second** instead is significantly more beneficial to automotive LiDAR systems. Relying on a higher point per second rate does nothing to signify the system's ability to accuracy and efficiency acquire objects of interest or hazards on the road. These types of systems have a hard time separating valuable information from the invaluable (or redundant) in a timely manner, causing latency and wasting energy. Therefore, it is more important to change the industry's way of thinking to favor the *quality* of the returned data over *quantity.*

Last, we have shown that it is critical to advance beyond measuring detection range, and instead, measure **object classification range** (i.e., how quickly an object can be identified and classified). It is not simply enough to quantify a distance at which a potential object can be detected at the sensor. One must also quantify the latency from the actual event to the sensor detection—plus the latency from the sensor detection to the CPU decision. Under this framework, the more attributes a LiDAR system can provide, the faster a perception system can classify.

While groundbreaking in their time, conventional LiDAR sensors passively search with scan patterns that are fixed in both time and space. A new generation of intelligent sensors moves from passive detection of objects to active search and identification of classification attributes of objects in real time. As perception technology and sensor systems evolve, it is imperative that metrics used to measure their capabilities also evolve.

The active iDAR system enables the type of toolkit that reduces latency and bandwidth in a dramatic way. It allows for the "Search, Acquire, and Act" functions to be implemented at the sensor level. Mimicking the process of human perception, new objects can be detected efficiently, classified with multiple supporting sensors, acted upon, fully interrogated for more information, or tracked with real-time data. Therefore, equipped with iDAR, an autonomous vehicle can spot hazards sooner and respond more quickly and accurately than it could using other sensor systems, avoiding accidents and giving credibility to the safety promise of self-driving vehicles.

With safety of paramount importance, these extended metrics should not only indicate what LiDAR systems are capable of achieving, but also how those capabilities bring vehicles closer to optimal safety conditions in real-world driving scenarios.

Throughout this series of white papers, AEye will continue to propose new, interconnected metrics that build on each other to help create a more complete and accurate picture of what makes a LiDAR system effective.

About the Authors

Blair LaCorte, CEO and Board Member

LaCorte is an accomplished leader and strategist with a long history of leveraging his change management skills to drive operational alignment and growth within companies. He has served as Global President of PRG, the world's largest live event technology and services company, CEO of XOJET, one of the fastest growing aviation companies in history and the largest private charter company in North America, and operating partner at TPG, a premier private equity firm with over $91B in global investments. LaCorte has also held numerous executive and general management positions in private and public technology companies including VerticalNet, Savi Technologies, Autodesk and Sun Microsystems. LaCorte graduated summa cum laude from the University of Maine and holds an MBA from Dartmouth's Tuck School of Business, where he later served as an executive fellow at the Center for Digital Strategies.

Luis Dussan, Founder, President, CTO and Board Member

Luis Dussan is a 20+ year veteran of the aerospace and defense industry. During his career at Lockheed Martin and Northrop Grumman, Dussan's research and development interests included smart and innovative sensor solutions for applications involving Information, Surveillance, Reconnaissance (ISR), Targeting, Fire Control, LADAR/LIDAR and Autonomy. He started his career at NASA working for the Jet Propulsion Lab in the Deep Space Network that communicated with NASA planetary and deep space probes. He then spent the bulk of his career at Lockheed Martin in their Missiles and Fire Control Division working in the Advanced Concepts group and developing state-of-the-art ISR and Targeting systems, such as the well-known LM Advanced Targeting Pod found in virtually every major US airborne military asset. After LM, Dussan went to Northrop Grumman Laser Systems and took the post of Chief Technologist where he was responsible for managing that division's Electro-Optical Sensors R&D. Dussan holds a BS in Electrical Engineering & Computer Science, an MS in Quantum Optics and an MS in Optics & Photonics. He put his PhD in Computational Physics on hold to start AEye.

Dr. Allan Steinhardt, Chief Scientist

An IEEE fellow, Dr. Allan Steinhardt is a sought-after expert on radar, missile defense, GMTI and space surveillance. He was Chief Scientist for DARPA, co-author of a book on adaptive radar, and Assistant Professor in Electrical Engineering and Applied Mathematics at Cornell University, where he performed funded research on sensor arrays and optimal detection capabilities. Dr. Steinhardt is a member of the National Academy of Sciences Naval Studies Board, and recipient of the US Defense Medal for Exceptional Public Service. He has also served as Chief Scientist at Booz Allen, the radar project lead at MIT Lincoln Laboratory and Director of Signal Processing for the defense industry with BAE/Alphatech.

Dr. Barry Behnken

Dr. Barry Behnken is a 25-year defense/aerospace executive with deep expertise in LiDAR and electro-optics, and extensive experience developing, managing and deploying disruptive imaging technologies for various applications throughout the world. He was Engineering Fellow for Emerging Laser Technologies at Raytheon Space and Airborne Systems, where he led the development and fielding of new DARPA LiDAR technologies. He served twenty years as a senior physicist and program manager in the US Air Force, including as Director of a $1.6B+ space satellite division, before retiring from active duty at the rank of lieutenant colonel. He holds BS, MS and PhD degrees in applied physics, and an MBA from the Haas School of Business at the University of California, Berkeley.

About AEye

AEye is the premier provider of high-performance, AI-driven LiDAR systems for vehicle autonomy, advanced driver-assistance systems (ADAS), and robotic vision applications. AEye's smart, software-configurable iDAR™ (Intelligent Detection and Ranging) platform combines solid-state, active LiDAR, an optionally fused low-light HD camera, and integrated deterministic artificial intelligence to capture more intelligent information with less data, enabling faster, more accurate, and more reliable perception. AEye has partnered with leading Tier 1s—such as Continental, Hella, and Aisin—and system integrators to configure and manufacture the sensor at scale to meet the diverse performance and functional requirements of autonomous and partially automated applications.