

How RTK Can Help Autonomous Vehicles See the World

Autonomous vehicles can't become fully reliable until they precisely know their location. Here's how real-time kinematic integration provides the accuracy needed to propel self-driving technology forward.

To varying degrees, autonomous vehicles (AVs) have been in development—and even in use—for decades. But technology has only recently begun to catch up to our imagination. Some experts now project that as many as 57% of passenger vehicles sold by 2035 could be [equipped](#) with true self-driving capabilities. Even where full autonomy won't yet be supported, cars will continue to come equipped with ever more [advanced driver-assistance systems](#) (ADAS).


And that's only [passenger vehicles](#). Google's autonomous DELIVERY vehicle fleet has already logged more than

300,000 miles. Waymo launched robotaxis, and Amazon is set to begin testing them in a few states. The first race of the Abu Dhabi Autonomous Racing League happened this year, and the Indy Autonomous Challenge has been active since 2021.


In short, momentum is shifting in favor of autonomous vehicles, and the commercial sector offers countless opportunities to leverage this fast-improving technology. Yet, just how quickly AVs become integral to everyday commercial operations hinges on their ability to see the world and react to it as well as or better than humans.

To do that, self-driving cars rely on a complex tech stack that processes information from an array of sensors (e.g., cameras), inertial measurement units (IMUs) (*Fig. 1*), Global Navigation Satellite System (GNSS) receivers, and [light detection and ranging \(LiDAR\)](#).

Each component of that sensor stack plays a crucial part in helping autonomous and ADAS-equipped vehicles navigate the world, but only one can provide a straightforward, convenient, absolute reference point for positioning: GNSS receivers. Still, traditional single-receiver positioning estimates are subject to significant errors of several meters or more, a dangerous



1. An IMU uses gyroscopes and accelerometers to determine angle and direction.



2. GNSS works with relative-positioning components of the sensor stack to provide important information for AVs and ADAS to make course corrections.

margin for moving vehicles. Achieving true precision—the kind that will allow more AVs to take the road safely—requires a reliable method for refining satellite-positioning estimates.

Real-time kinematic (RTK) corrections can bring satellite-positioning estimates to within centimeter-level accuracy, providing a far more reliable source of absolute positioning for AVs. To understand how, we must delve into the details of the self-driving stack, explore the role of satellite data within that stack, and see what RTK can add to the equation, especially when properly integrated via sensor fusion.

How Self-Driving Cars “See” the World

For AVs to safely navigate the world, they must accurately “see” everything happening in their environment. This “seeing” represents the first three tasks of the self-driving stack: sensing, perception, and localization. The car must sense objects and traffic markers in and around its path, correctly perceive what those objects are and what they’re doing, and then accurately locate itself in relation to those objects. Only then can it move to the other two tasks of planning and control.

In this article, we’ll concern ourselves primarily with the “seeing” tasks of the tech stack, as we’re dealing first and foremost with how AVs can more accurately determine their location. As it turns out, seeing and perceiving our surroundings—and our place in those surroundings—is something we humans take for granted. AVs must rely on a whole suite of tools to accomplish what we do subconsciously with just our eyes and brain.

These complex tools include:

- **Cameras:** Front, rear, and side visible cameras provide an immediate 360-degree picture of the world, but they’re limited by light and weather conditions. Cameras

built with infrared sensors, night vision, or other sensors are designed to overcome some of these limitations.

- **Radar sensors:** Pulsating radio waves provide information about an object’s speed and distance, but typically their resolution is lower and data sparser than that of other technologies.
- **LiDAR sensors:** LiDAR uses laser pulses to create detailed, accurate pictures of a vehicle’s surroundings. While it’s possible to create maps for positioning [with LiDAR alone](#), this is a costly, time-consuming proposition.
- **Ultrasonic sensors:** Similar to how bats use echolocation, cars can use short-range ultrasound to locate nearby moving or stationary objects. These sensors are most helpful for slow, close-range activities like parking or backing up in a tight space.
- **Inertial measurement units (IMUs):** IMUs use gyroscopes, accelerometers, and (in some cases) magnetometers to measure the specific force and angular rate of an object. This technology provides the AV with a sort of internal sense of direction.
- **GNSS receivers:** GNSS like GPS, GLONASS, GALILEO, and Bei-Dou transmit satellite location and timing information to receivers, which the receiver can use to calculate the absolute positioning of the AV within space. As we’ll explore further, a variety of environmental factors can limit the accuracy of these positioning estimates.

Together, these various components of the tech stack provide a wealth of information. But, without a process for integrating and evaluating this information, we would be no closer to creating an autonomous vehicle. Moving from sensing to perception, localization, and beyond relies on a complex mixture of computer vision and sensor fusion.

Many of the high-profile accidents involving AVs have occurred due to some shortcomings in this process of integration and decision-making. Nonetheless, the more accurate the inputs are from each component of the tech stack, the better the starting point for computer vision and sensor fusion to do their jobs.

In the past decade, AV designers have relied on building high-definition (HD) maps that require survey-grade data processed by simultaneous localization and mapping (SLAM) or SLAM-like algorithms. However, this method comes with a high computational overhead. And, where loop closure isn’t possible, some degree of drift is inevitable. In places like highways, where scenery changes frequently, maintaining the necessary detail of an HD map can be incredibly expensive.

Ideally, you would have a reliable, affordable source of authority on which to base the sensor-fusion process. And that's exactly what you find in GNSS with RTK.

The Role and Limitations of GNSS in the Autonomous Vehicle Tech Stack

As the source of absolute positioning within the AV tech stack, GNSS receivers (Fig. 2) play a critical role in making autonomous driving possible. The basic design is simple: These receivers capture signals from satellite constellations, which contain information about where each satellite is located and what time the signal was sent. As long as the receiver is within view of at least four orbiting satellites, it can translate this information into a fairly accurate estimate of the receiver's location in space.

At the most basic level, this is what's happening on every person's GPS-enabled smartphone when they use a navigation app (although phone location typically also incorporates data from the phone's IMU, Wi-Fi signals, and cellular network, and navigation snaps to roads). And it works quite well when you have a human driver involved to account for any errors.

In the case of self-driving cars, the limitations of real-time, single-receiver GNSS positioning become glaringly obvious. In the above scenario, accurate positioning estimates rarely occur because GNSS signals are subject to all sorts of interference (Fig. 3).

They include:

- **Satellite and receiver clock errors:** These clocks are biased and drift slowly from absolute accuracy, introducing small but significant errors if they're not correctly modeled.

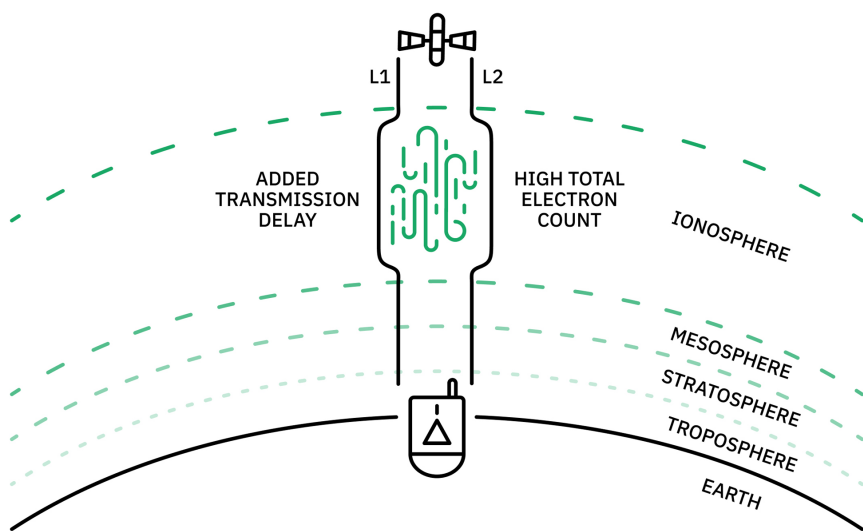
- **Multipath and non-line-of-sight reception:** GNSS readings are most accurate where there's a direct line of sight between the satellite and the receiver—in the open plains or desert, for instance. Introduce obstructions like trees or tall buildings, and you'll get signal delays or multipath disturbances from signals bouncing off these objects. That's why your smartphone's GPS may struggle in crowded urban areas when you're surrounded by skyscrapers.
- **Ephemeris errors:** The ephemeris is a set of satellite orbit parameters that allow for the calculation of satellite coordinates at various points in time. Factors like the moon's gravitational pull or solar radiation can all affect satellite positioning, and the model used in real-time GNSS for calculating coordinates from the ephemeris doesn't account for all of these variables. This leads to inaccurate estimates of satellite positioning. If the satellite position is wrong, the receiver's estimate of its own location will likewise be inaccurate.
- **Ionospheric and tropospheric delays:** These two layers of the Earth's atmosphere introduce the most significant delays and distortions to satellite signals on their way to the surface. Of the two, ionospheric errors are more significant, sometimes causing receiver positioning errors of several meters or more.

In many applications, these delays are insignificant enough to be ignored. Ultimately, though, they can add up to errors in positioning estimates of several meters or more. It's not difficult to imagine how that could be disastrous for self-driving cars.


In the best-case scenario, a driverless delivery car drops your package on your neighbor's lawn instead of right in front of your house. In much worse cases, you have self-driving cars causing fatal accidents based on erroneous positioning information.

Of course, numerous failsafes and corrections are built into the tech stack, and the relative positioning information provided by an AV's many sensors is invaluable for preventing serious mistakes.

However, each part of the stack has inherent limitations. Standard visible-spectrum cameras can't handle certain weather or lighting conditions, while radar comes up short for precise distinctions between objects. Adding a reliable source of absolute positioning information helps diminish the significance of these shortcomings and contributes to a more accurate sensor-fusion process.



3. Ionospheric disturbances are one of the most common sources of GNSS signal delays, resulting in positioning errors on the ground.



4. Placing a GNSS receiver close enough to another device with a known, fixed position can cancel out most sources of error to arrive at an accurate position estimate for your receiver.

Yet, merely integrating GNSS into the stack—subject to these various error sources—may not be enough. That’s why successful wide-scale commercial deployment demands an absolute method for securing more accurate, reliable positioning data.

Locking in on Location with Real-Time Kinematic Corrections

RTK positioning (*Fig. 4*) is based on a simple premise: If you can place a GNSS receiver close enough to another

device that has a known, fixed position, you can effectively cancel out most sources of error to arrive at an accurate position estimate for your receiver.

It’s astoundingly simple but extraordinarily effective. By setting up a base station within an appropriate range—usually 50 kilometers or less—atmospheric delays, clock and ephemeris inaccuracies, and even some multipath errors simply drop out. As a result, RTK gives you GNSS positioning estimates that are accurate to within a few centimeters.

Now, it’s possible to implement RTK positioning with your own base station and an RTK-enabled receiver. However, this DIY setup is often insufficient for robust commercial operations, especially where AVs are involved. For instance:

- DIY base stations may be less accurate due to setup mistakes or poor practices in establishing base locations.
- Homemade stations are often built with lower-quality antennas and receivers.
- A single base station offers a limited coverage area for an AV, and building multiple base stations can get quite expensive and time-consuming.

These limitations are prohibitive for meeting the high-precision demands of self-driving vehicles. In most cases, it’s



5. Faktion relies on RTK to help its self-driving delivery vehicles achieve centimeter-level precision.

better to rely on an established network of RTK base stations to correct GNSS estimates. This removes the uncertainties associated with DIY and, assuming a large enough network, enables AVs to move freely across a wide range without fear of losing accuracy.

Using RTK and INS to Support Autonomous Navigation

Even with RTK, GNSS receivers can't always estimate absolute position well. Tunnels completely obscure satellite signals, for example, and dense cityscapes may create more multipath errors than RTK can cancel out. Fortunately, there's a way to integrate GNSS and RTK into the self-driving stack to localize accurately and update maps in real-time.

By leveraging an inertial navigation system (INS) to "fill in the gaps" of GNSS measurements, you can combine absolute and relative positioning to provide a reliable source of geo-registration for an AV's map. Essentially, this combines the highly precise data from an RTK-enabled GNSS receiver with the directional information from an IMU to create a highly accurate picture of the vehicle's location in the real world.

From here, we can begin to see the full value of integrating INS with RTK into the entire process of HD mapping and localization. By integrating RTK-INS as priors for your SLAM process, the "first guess" for optimization is already very accurate.

As a result, you can minimize much of the computational overhead associated with map acquisition and limit drift where loop closure isn't possible. Similarly, where scenery changes frequently, INS with RTK can reduce the level of detail that must be stored in the map, or you can even use it on its own for localization.

This approach could be used with newer methods like deep learning, which may even take the place of HD mapping in some cases. However, RTK-INS can still enhance accuracy in this method, offering a low-cost, convenient way to both obtain better ground truth for model training and constrain prediction.

Real-Time Kinematic Driving in the Real World

This is more than a mere pipe dream. Such a level of sensor fusion is already possible, and it's being used by real vehicles in the real world. Consider the following examples.

[Faction](#) (Fig. 5) builds systems to enable driverless operation for last-mile curbside delivery and logistics fleets. The company recently deployed Point One's Atlas INS for precise, RTK-corrected positioning data combined with real-time relative positioning from the built-in Atlas IMU. Faction is already testing this system on small electric vehicles for last-mile deliveries within busy city centers and suburbs.

Paving the Way for the Self-Driving Cars of the Future

Autonomous-vehicle technology is accelerating quickly, and self-driving cars aren't far from becoming an everyday reality. If we're to make that transition safely, it will take a highly integrated sensor stack with consistently reliable absolute and relative positioning information.

By refining GNSS positioning estimates with RTK, we can provide a stronger absolute reference point for AVs. When combined with IMU data and other relative sensor information through powerful sensor fusion, this technology enables AVs to not merely see the world as we do, they get a more complete, accurate picture than any human could ever dream of.

From there, the only remaining task is to ensure they're smart enough to respond to that picture better than a human driver.

Karen Leineweber currently leads the Point One Navigation computer vision team, and has held a key role in building the company's satellite navigation and sensor fusion systems. Prior to joining Point One, she was an Engineer in the space robotics division at MDA Corporation, where she was part of the core team responsible for the advancement of Canadian planetary rover technology, particularly in the navigation, control, and visual odometry segments.

Karen received her PhD in Aerospace Engineering from Cornell University. Her graduate work consisted of developing estimation methods that use GNSS/GPS signal data to support atmospheric and space weather research. Karen also has a BS from Columbia University in Applied Physics.