# Geometry-Based Discrimination of Authentic and Spoofed GNSS Signals Using Double Differencing

Jade Babcock-Chi, Todd Walter, *Stanford University*

## BIOGRAPHY

**Jade Babcock-Chi** is a PhD student in the Department of Aeronautics and Astronautics at Stanford University working under the guidance of Professor Todd Walter in the GPS Lab. Prior to joining the lab, Jade received her B.S. and M.S. in Aerospace Engineering at the University of Colorado Boulder. Her research interests include GNSS integrity and navigation in challenging environments.

**Todd Walter** received his Ph.D. in Applied Physics from Stanford University in 1993. He is a Research Professor in the Department of Aeronautics and Astronautics at Stanford University. His research focuses on implementing high-integrity air navigation systems. He has received the ION Thurlow and Kepler awards. He is also a fellow of the ION and has served as its president.

## ABSTRACT

GNSS spoofing and meaconing pose a growing threat to safety-critical systems reliant on GNSS for positioning, navigation, and timing. This work presents a pseudorange double-difference framework for discriminating between authentic and spoofed GNSS signals across a network of spatially separated receivers, and demonstrates its application to spoofer localization via time-difference-of-arrival (TDOA). Three $C/N_0$-based noise models are empirically derived using a code-minus-carrier observable, and a geometric separation condition is enforced to ensure reliable hypothesis discrimination between all-authentic and all-spoofed signal quartets. The methodology is validated against three spoofing scenarios from Jammertest 2025 using GPS L1 C/A and Galileo E5a measurements across a network of ten receivers, successfully localizing a close-proximity spoofer to within a few meters and recovering directionally consistent transmitter estimates for more geometrically challenging configurations. Limitations and directions for future work are discussed, with emphasis on a network-wide joint classification and localization framework to build upon the current quartet-based enumeration.

## I. INTRODUCTION

The Global Navigation Satellite System (GNSS) underpins critical positioning, navigation, and timing (PNT) functions across transportation, power systems, telecommunications, and emergency services, making it an attractive target for intentional interference. Unlike jamming, which causes outright denial of service, spoofing attacks inject counterfeit signals that capture a receiver's tracking loops and steer the navigation solution while preserving apparently nominal indicators such as signal-to-noise ratio and dilution of precision (Psiaki & Humphreys, 2016; Tippenhauer et al., 2011). As GNSS becomes more deeply embedded in safety-critical infrastructure, there is a growing need for methods that can reliably separate authentic and spoofed signals using modest additional hardware, and that remain effective under dynamic and increasingly sophisticated attack strategies (Jafarnia-Jahromi et al., 2012; Qin et al., 2018).

Pseudorange double differencing is a well-established technique in high-precision GNSS that cancels common-mode errors — satellite clock offsets, receiver clock bias, and a substantial portion of atmospheric delays — by differencing simultaneously across receivers and satellites (Hofmann-Wellenhof et al., 2008; Misra & Enge, 2012). This structure naturally extends to spoofing detection: signals from a single spoofing transmitter share a common spatial origin and induce correlated observables across the receiver network, whereas authentic signals arrive from widely separated satellites with distinct lines of sight (Motella et al., 2011; Zhu et al., 2019). In the double-difference domain, these two signal classes produce characteristically different residual signatures that can be exploited as a separation metric (Qin et al., 2018; Wang et al., 2017).

Existing multi-receiver methods, including the carrier-phase graph-based approach of Jafarnia-Jahromi et al. (2016), generally rely on temporal averaging or long integration windows to stabilize decision statistics, and often require a minimum number of spoofed channels before separation is feasible, limiting responsiveness to rapidly evolving or mobile spoofing threats. In this work, we instead pursue an epoch-by-epoch signal separation strategy in the double-difference pseudorange domain. At each epoch, double differences formed across carefully selected receiver–satellite quartets are used to classify signals as authentic or spoofed without requiring time averaging or a minimum number of compromised channels. By operating on measurement snapshots, the method remains effective even when the spoofing transmitter is in motion. We presently restrict attention to a single spoofing transmit antenna and focus on quartets of four measurements as a tractable starting point, while laying the

groundwork for a scalable network-wide signal separation framework capable of continuously identifying authentic and spoofed signals in real time.

## II. METHODOLOGY

This section describes the measurement models, test statistic construction, noise assumptions, geometry-based separability metric, decision logic, and validation methodology used to discriminate between authentic and spoofed GNSS signals using pairs of receivers and satellites.

### 1. Pseudorange Measurement Models

We consider a scenario in which two spatially separated receivers observe signals from multiple satellites, where the received signals may be either authentic or spoofed. Separate pseudorange measurement models are defined for these two cases.

*a) Spoofed pseudorange model*

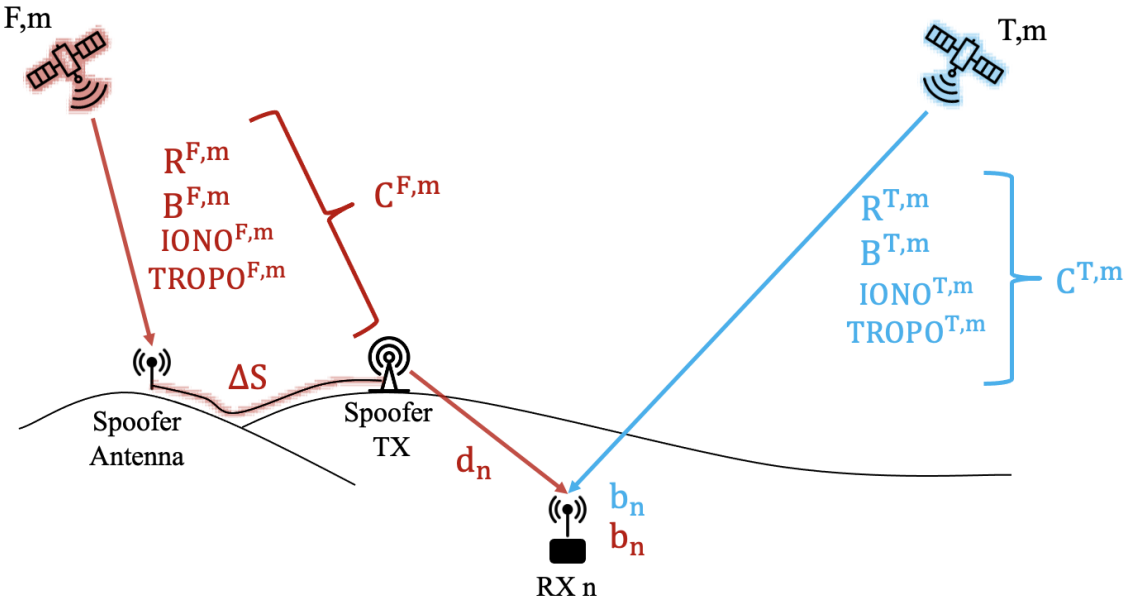


**Figure 1:** Pseudorange double difference scenario for spoofing or meaconing

A meaconer receives GNSS signals and retransmits signals toward victim receivers. A spoofer generates synthetic GNSS signals and broadcasts those signals toward victim receivers. Our model assumes spoofing and meaconing are synonymous. For receiver $n$ observing a spoofed signal corresponding to satellite $m$, the spoofed pseudorange measurement is modeled as

$$\rho_n^{F,m} = R^{F,m} + B_{sat}^{F,m} + \mathrm{IONO}^{F,m} + \mathrm{TROPO}^{F,m} + \Delta S + d_n + b_n + \epsilon_n^{F,m}. \tag{1}$$

Here, $R^{F,m}$ represents the geometric range between the satellite signal source and the transmitting antenna. The terms $B_{sat}^{F,m}$, $\mathrm{IONO}^{F,m}$, and $\mathrm{TROPO}^{F,m}$ account for satellite clock bias and propagation delays incurred by the transmitter, as seen in Fig 1. The term $\Delta S$ represents additional delays introduced between the spoofer receiver and transmitter, including cable delays and processing latency. The distance $d_n$ denotes the geometric distance from the spoofer transmitter to receiver $n$, while $b_n$ is the receiver clock bias. The term $\epsilon_n^{F,m}$ represents the measurement noise.

For the purposes of hypothesis discrimination, the absolute values of the satellite-dependent and spoofer-dependent delays are not of interest. These terms are therefore grouped into a satellite-specific nuisance parameter

$$C^{F,m} = R^{F,m} + B_{sat}^{F,m} + \mathrm{IONO}^{F,m} + \mathrm{TROPO}^{F,m} + \Delta S. \tag{2}$$

The spoofed pseudorange model used in simulation is thus expressed as

$$\rho_n^{F,m} = C^{F,m} + d_n + b_n + \epsilon_n^{F,m}. \tag{3}$$

*b) Authentic pseudorange model*

For authentic GNSS signals, the pseudorange measured by receiver $n$ from satellite $m$ is modeled as

$$\rho_n^{T,m} = R_n^{T,m} + B_{sat}^{T,m} + \mathrm{IONO}^{T,m} + \mathrm{TROPO}^{T,m} + b_n + \epsilon_n^{T,m}. \tag{4}$$

Here, $R_n^{T,m}$ is the geometric range between satellite $m$ and receiver $n$. Satellite-dependent propagation delays and clock biases are grouped into a nuisance parameter

$$C^{T,m} = B_{sat}^{T,m} + \mathrm{IONO}^{T,m} + \mathrm{TROPO}^{T,m}. \tag{5}$$

The authentic pseudorange model used in simulation becomes

$$\rho_n^{T,m} = R_n^{T,m} + C^{T,m} + b_n + \epsilon_n^{T,m}. \tag{6}$$

The receiver tracks only one signal per satellite. And for now, we assume that any given measurement can only take on a binary value of spoofed or authentic. The possibility of a bad measurement is entirely possible but is not considered in this paper.

## III. TEST STATISTIC, GEOMETRIC FACTOR, AND ERROR DISTRIBUTION CHARACTERIZATION

This section describes the test statistic construction, noise assumptions, geometry-based separability metric, decision logic, and data processing methodology used to discriminate between authentic and spoofed GNSS signals using pairs of receivers and satellites.

### 1. Pseudorange Double Difference Test Statistic

Consider two receivers, indexed 1 and 2, each observing two satellites, labeled $A$ and $B$. Let superscripts $T$ and $F$ denote authentic (true) and spoofed (fake) signals, respectively. The pseudorange double difference test statistic is defined as

$$x = \left(\rho_1^A - \rho_2^A\right) - \left(\rho_1^B - \rho_2^B\right), \tag{7}$$

where $\rho_n^m$ denotes the pseudorange measurement at receiver $n$ from satellite (or spoofed signal corresponding to satellite) $m \in \{A,\, B\}$. The expected value of $x$ depends on the hypothesis under test.

*a) All-authentic hypothesis*

When all four pseudoranges are authentic, substituting the authentic model gives

$$\begin{aligned}
\rho_1^A &= R_1^{T,A} + C^{T,A} + b_1 + \epsilon_1^{T,A},\\
\rho_2^A &= R_2^{T,A} + C^{T,A} + b_2 + \epsilon_2^{T,A},\\
\rho_1^B &= R_1^{T,B} + C^{T,B} + b_1 + \epsilon_1^{T,B},\\
\rho_2^B &= R_2^{T,B} + C^{T,B} + b_2 + \epsilon_2^{T,B}.
\end{aligned} \tag{8}$$

The satellite-dependent nuisance parameters $C^{T,m}$ and receiver clock biases $b_n$ cancel in the double difference, yielding

$$x = \left(R_1^{T,A} - R_2^{T,A}\right) - \left(R_1^{T,B} - R_2^{T,B}\right) + \nabla\Delta\varepsilon, \tag{9}$$

where $\nabla\Delta\varepsilon$ is the double difference of the individual pseudorange noise terms. Defining the *geometric factor*

$$L = \left(R_1^{T,A} - R_2^{T,A}\right) + \left(R_1^{T,B} - R_2^{T,B}\right), \tag{10}$$

the expected value of the test statistic under the all-authentic hypothesis is $x \approx L + \nabla\Delta\varepsilon$, or equivalently $x - L \approx 0$.

*b) All-spoofed hypothesis*

When all four pseudoranges are spoofed, substituting the spoofed model gives

$$\begin{aligned}\rho_1^A &= C^{F,A} + d_1 + b_1 + \epsilon_1^{F,A},\\ \rho_2^A &= C^{F,A} + d_2 + b_2 + \epsilon_2^{F,A},\\ \rho_1^B &= C^{F,B} + d_1 + b_1 + \epsilon_1^{F,B},\\ \rho_2^B &= C^{F,B} + d_2 + b_2 + \epsilon_2^{F,B}.\end{aligned} \tag{11}$$

Again, the satellite-dependent and receiver clock terms cancel, and because the spoofer transmitter-to-receiver distances $d_n$ are common to both satellites, the double difference evaluates to

$$x = \nabla\Delta\varepsilon \approx 0. \tag{12}$$

The two hypotheses therefore produce distinct expected values for the test statistic:

$$\mathcal{H}_T: \quad x \approx L \quad \text{(all authentic)}, \tag{13}$$

$$\mathcal{H}_F: \quad x \approx 0 \quad \text{(all spoofed)}. \tag{14}$$

The geometric factor $L$ therefore acts as the separation between the two hypothesis distributions, and its magnitude directly governs the ability to discriminate between the two cases. When the authentic satellite geometry is poor (i.e. when the geometric ranges from the two satellites to the two receivers are nearly equal) $L$ approaches zero and the two distributions overlap, potentially leading to false detections. This is discussed further in Section III.3.

### 2. Noise Model and Error Distribution

The double difference noise term $\nabla\Delta\varepsilon$ is the sum of four independent pseudorange noise contributions:

$$\nabla\Delta\varepsilon = \epsilon_1^A - \epsilon_2^A - \epsilon_1^B + \epsilon_2^B. \tag{15}$$

Each individual pseudorange error $\epsilon_n^m$ is modeled as a zero-mean Gaussian random variable with standard deviation $\sigma_{\rho_n^m}$:

$$\epsilon_n^m \sim \mathcal{N}\left(0,\, \sigma^2_{\rho_n^m}\right). \tag{16}$$

Under the assumption of independence, the standard deviation of the test statistic is therefore

$$\sigma_x = \sqrt{\sigma^2_{\rho_1^A} + \sigma^2_{\rho_2^A} + \sigma^2_{\rho_1^B} + \sigma^2_{\rho_2^B}}. \tag{17}$$

The distribution of the test statistic under each hypothesis is thus

$$\mathcal{H}_T: \quad x \sim \mathcal{N}(L,\, \sigma_x^2), \tag{18}$$

$$\mathcal{H}_F: \quad x \sim \mathcal{N}(0,\, \sigma_x^2). \tag{19}$$

A key assumption made here is that the noise distribution is the same for both authentic and spoofed signals. This is justified by the observation that the spoofer retransmits signals whose received carrier-to-noise density ratio ($C/N_0$) is similar to that of authentic satellites. Since $C/N_0$ is the primary driver of pseudorange noise, the noise standard deviation $\sigma_\rho$ is modeled as a function of $C/N_0$ alone, rather than as a function of elevation angle. Elevation angle is not used because the elevation angle of the spoofing transmitter is not known a priori. Given that $C/N_0$ measurements are available for both spoofed and authentic signals, $\sigma_\rho$ is empirically characterized as a function of $C/N_0$ and the resulting curve is used to assign noise variances to individual pseudoranges. The characterization of $\sigma_\rho(C/N_0)$ is presented in Section IV.

It must be noted that the zero-mean Gaussian assumption is a simplification. In practice, GNSS pseudorange errors are subject to multipath, non-line-of-sight reception, and other environmental effects that can introduce systematic biases, rendering the errors non-zero-mean. These effects are not modeled here and represent a known limitation of the methodology that warrants further investigation.

## 3. Decision Framework

To classify a quartet of measurements, the following decision logic is applied at each epoch based on our single test statistic $x$. Recalling that the expected value of $x$ is approximately $L$ under the all-authentic hypothesis and approximately zero under the all-spoofed hypothesis, a quartet is declared *all-authentic* if $|x - L| < 2\sigma_x$, *all-spoofed* if $|x| < 2\sigma_x$, and *inconclusive* otherwise. The threshold of $2\sigma_x$ was chosen heuristically: it is large enough to avoid discarding too many quartets due to non-Gaussian heavy tails in the noise distribution, which may arise from multipath and other real-world effects, while remaining small enough to maintain discriminative power between the two hypotheses.

For the two decision regions to be non-overlapping, the hypothesis means must be sufficiently separated relative to the noise level. Requiring that neither $2\sigma_x$ acceptance window extend past the midpoint between the two means yields the condition

$$|L| \geq 4\sigma_x. \tag{20}$$

When this condition is not met, the distributions of $x$ under $\mathcal{H}_F$ and $\mathcal{H}_T$ overlap within the decision boundaries, and the quartet is marked inconclusive regardless of the observed value of $x$. This situation arises when the authentic satellite geometry is poor (i.e when the two satellites are nearly equidistant from both receivers) causing $L$ to approach zero and the two hypotheses to become indistinguishable from noise alone.

### *a) Additional hypotheses and false-detection risk*

For a quartet of four pseudorange measurements, each of which may independently be either authentic or spoofed, there are $2^4 = 16$ possible hypotheses in total. The two hypotheses considered above — all-authentic and all-spoofed — are only two of these sixteen. The remaining fourteen represent mixed cases such as three authentic and one spoofed, or two of each. These mixed hypotheses are not directly testable with the present framework, as their distributions under the test statistic $x$ depend on the receiver geometry, the spoofer transmitter location, the spoofed satellite geometry, and the spoofer-induced delays. These factors combine in ways that are difficult to model without additional a priori knowledge.

In practice, it is expected that mixed hypotheses produce values of $x$ that are large in magnitude and therefore do not satisfy either of the $2\sigma_x$ thresholds. This reduces the probability of a false positive classification. However, this assumption cannot be guaranteed for all geometries and spoofing configurations, and the sensitivity of the detection framework to mixed hypotheses represents a direction for future work.

## 4. Data Processing Methodology

The following processing pipeline is applied to the collected data.

### *a) Receiver pairing and satellite selection*

All possible pairs of receivers within the network are enumerated. For each receiver pair, the set of satellites whose pseudoranges are simultaneously observed by both receivers is identified. This intersection forms the pool from which satellite quartets are drawn for double-difference testing.

### *b) Epoch-by-epoch classification*

For each receiver pair, the data are processed at 1 Hz. At each epoch, all valid quartets of two satellites and two receivers are tested using the decision framework described in Section III.3. Each quartet is labeled as all-authentic, all-spoofed, or inconclusive. An individual satellite signal is deemed authentic if at least one quartet in which it participates yields an all-authentic classification, and spoofed if at least one quartet yields an all-spoofed classification. The current implementation evaluates all combinations exhaustively, so a given signal may be reclassified multiple times across different quartets. Only a single classification is required to assign a label; however, more efficient selection strategies will be explored in future work.

### *c) Transmitter Localization and Validation*

Once at least one authentic and one spoofed signal per receiver have been identified at a given epoch, distance-difference measurements to the spoofer transmitter are computed following the methodology of Babcock-Chi et al. (2025). This is repeated over all valid authentic–spoofed signal pairings and receiver pairs, and the mean distance difference at each epoch is used as input to a two-dimensional nonlinear least-squares (NLLS) solver for TDOA-based transmitter localization. Successful convergence of the NLLS solution serves as an indirect validation of the signal classification: if authentic and spoofed signals had been incorrectly separated, the resulting TDOA measurements would be mutually inconsistent and the solver would fail to produce an accurate position estimate.

*d) Limitations*

A fundamental requirement of this methodology is the simultaneous presence of at least two authentic and two spoofed signals in the observed signal set. If only one spoofed satellite signal is present, no valid all-spoofed quartet can be formed and the classification framework cannot be applied to that measurement. All other authentic measurements can still be classified. Incorporating additional GNSS frequencies and constellations increases the number of available measurements improves both the probability of forming valid quartets and the robustness of the classification through redundancy.

## IV. RESULTS: STANFORD ROOFTOP

*a) Pseudorange Noise Characterization via Code-Minus-Carrier*

Pseudorange noise was characterized empirically using data collected from two u-blox F9P receivers mounted on the roof of the Durand Building at Stanford University as seen in Fig. 2a. A code-minus-carrier (CMC) observable was computed for each signal to isolate the pseudorange noise, removing the common-mode contributions of geometry, clock, and atmospheric delays. The CMC values from data from both receivers were grouped into integer $C/N_0$ bins (in dB-Hz) as seen in Fig. 2b, and the standard deviation within each bin was taken as the empirical noise estimate $\hat{\sigma}_\rho$ at that $C/N_0$ level, yielding the set of data points shown in Fig. 2c.

Three models were fit to these empirical estimates to enable extrapolation down to $C/N_0$ values as low as 10 dB-Hz, which may be encountered for weak or attenuated signals as seen in Fig. 2c. The first two models are polynomial fits of second and third order respectively, obtained via least-squares regression on the binned data points. The third model follows one of the standard parametric forms (Medina et al. 2018),

$$\sigma_\rho(C/N_0) = a + b \cdot 10^{-(C/N_0\,/\,10)}, \tag{21}$$

where $a$ and $b$ are scalar coefficients fit to the empirical data. This functional form is physically motivated, as thermal noise power scales inversely with the linear signal-to-noise ratio. Together, the three models provide a means of assigning a noise variance $\sigma^2_{\rho_n^m}$ to each pseudorange measurement based solely on its observed $C/N_0$. The validity of each model is assessed using independent data sets collected during Jammertest 2025, as described in Section V.

## V. RESULTS: JAMMERTEST 2025, 09/16/25

Experimental validation was performed using data collected during Jammertest 2025 across a network of ten receivers, whose layout was identical across all three scenarios tested. GPS L1 C/A and Galileo E5a were the only signals simultaneously present across a sufficient number of receivers with a mix of spoofed and authentic measurements (a requirement for both the classification framework and TDOA localization) as all other constellations and frequencies were either fully degraded or fully spoofed during the test periods. For each scenario, the active receiver subset was selected based on Galileo E5a availability and proximity to the transmitter under test; incorporating additional constellations and frequencies is left as future work (Section VII).

For each scenario, $C/N_0$ time histories for GPS L1 C/A (shown in dark red) and Galileo E5a (shown in yellow) are presented as a diagnostic tool (see Fig. 3-5, 9-12, 16-18). Signal strength directly governs pseudorange noise through the $\sigma_\rho(C/N_0)$ models developed in Section IV.0 a), and therefore directly affects localization accuracy. Low $C/N_0$ values (particularly those requiring extrapolation beyond the range of the empirical training data) introduce additional uncertainty into the noise characterization and, by extension, into the quartet classification and distance-difference estimates. Three spoofing scenarios are considered, corresponding to the Mountain, Community, and Cemetery transmitter locations, which represent a static meaconer, a close-proximity spoofer, and a more distant spoofer, respectively. All hyperbolic and TDOA localization plots presented are computed using the second-order polynomial fit of the $C/N_0$–$\sigma_\rho$ characterization described in Section IV.0 a).

### 1. Mountain Transmitter (Meaconer)

The Mountain transmitter operated as a meaconer, retransmitting authentic GNSS signals from an elevated static location outside the receiver network. As shown in Fig. 6, a cluster of position estimates is visible near $69°16'45''$N, $16°01'$E, corresponding to the known Mountain transmitter location and the spoofed position reported by one of the captured receivers. The $C/N_0$ measurements for this scenario are consistent across the test period as seen in Fig. 3 - 5, with no large outliers, indicating stable signal conditions and reliable noise characterization. The TDOA position estimates shown in yellow, however, exhibit a clear and systematic bias away from the known transmitter location as seen in Fig. 8. This bias is not accounted for in the current measurement model and is not attributable to noise alone given the consistency of the individual estimates. Understanding the source of this bias (which may relate to unmodeled multipath at the meaconer, asymmetric propagation conditions, or residual systematic errors in the distance-difference computation) is an ongoing area of research. Despite the inability to recover a precise point solution, the hyperbolic lines of position as seen in Fig. 7 are broadly consistent with the correct direction of arrival from


**(a)** Stanford rooftop receiver locations.

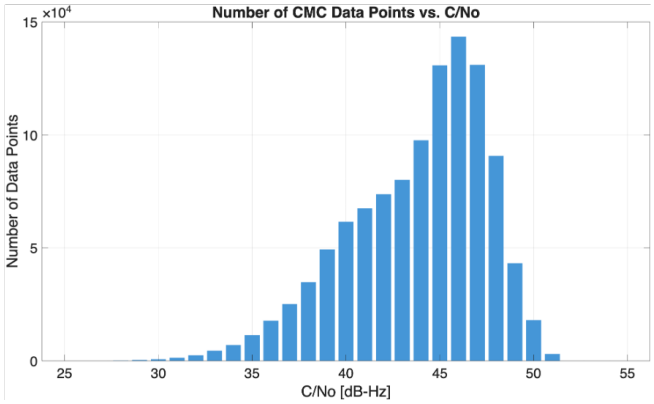


**(b)** Binned CMC data by C/No value.

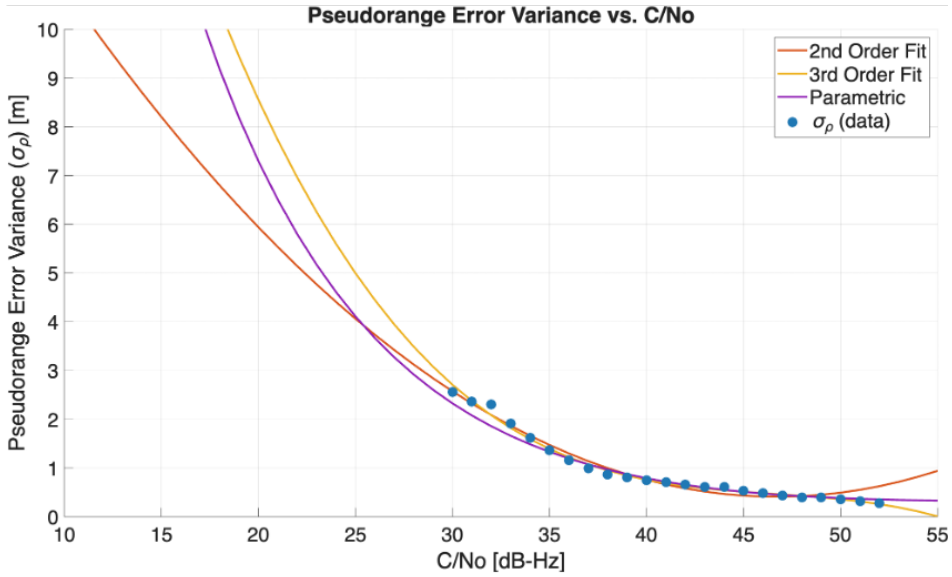


**(c)** Standard deviation of CMC bin data with 2nd order, 3rd order, and parametric fits.

**Figure 2:** Rooftop analysis showing measurement plots and model fit results.

the receiver network to the transmitter, suggesting that a direction-of-arrival estimate remains meaningful even with unfavorable geometry and receiver biases.

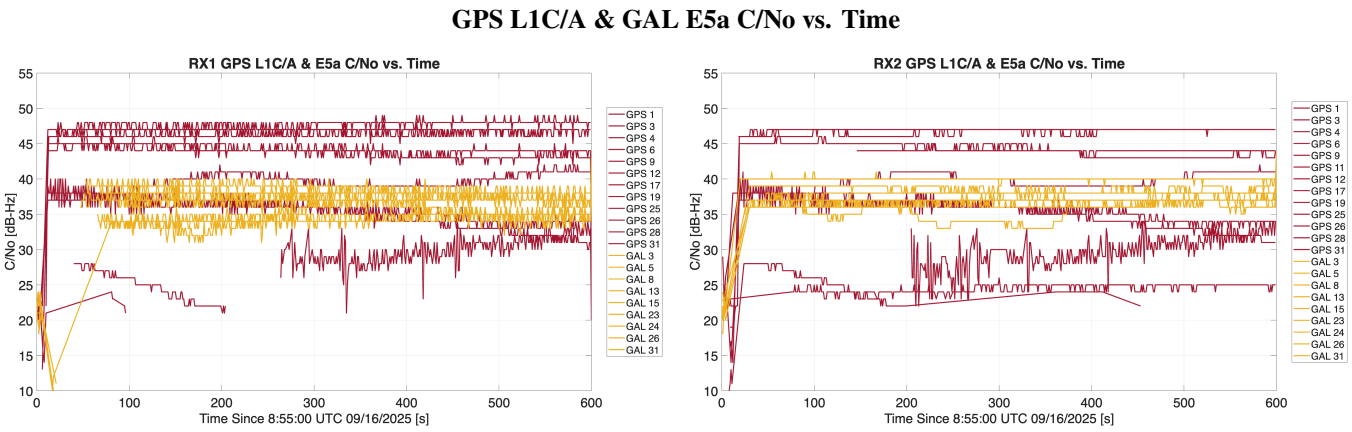


**Figure 3:** RX1 C/No during spoofing from the mountain.

**Figure 4:** RX2 C/No during spoofing from the mountain.

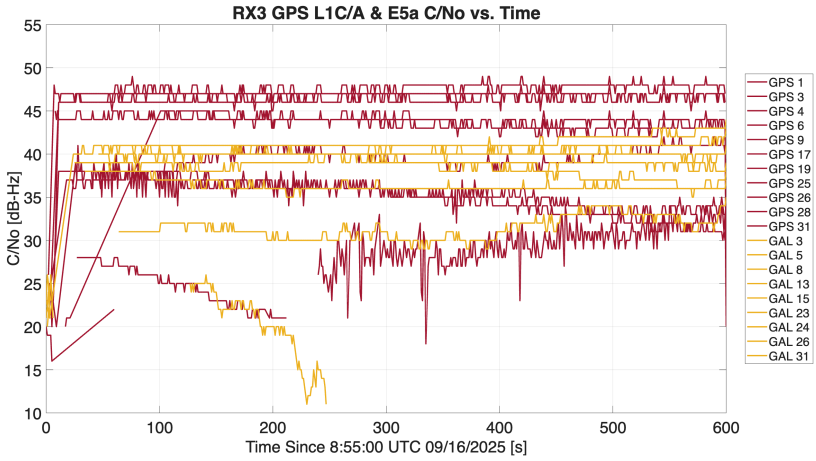


**Figure 5:** RX3 C/No during spoofing from the mountain.

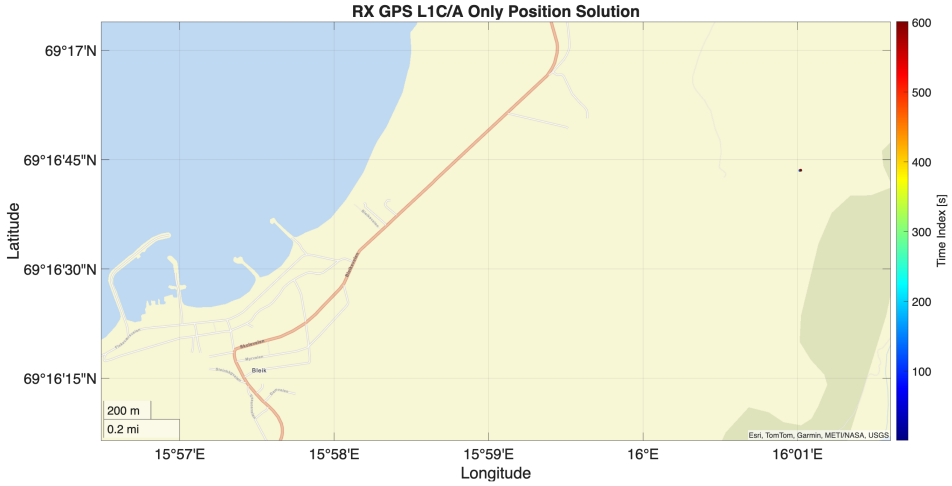


**Figure 6:** Example of full RX capture from mountain with GPS L1C/A measurements only as a function of time (stationary meaconing).

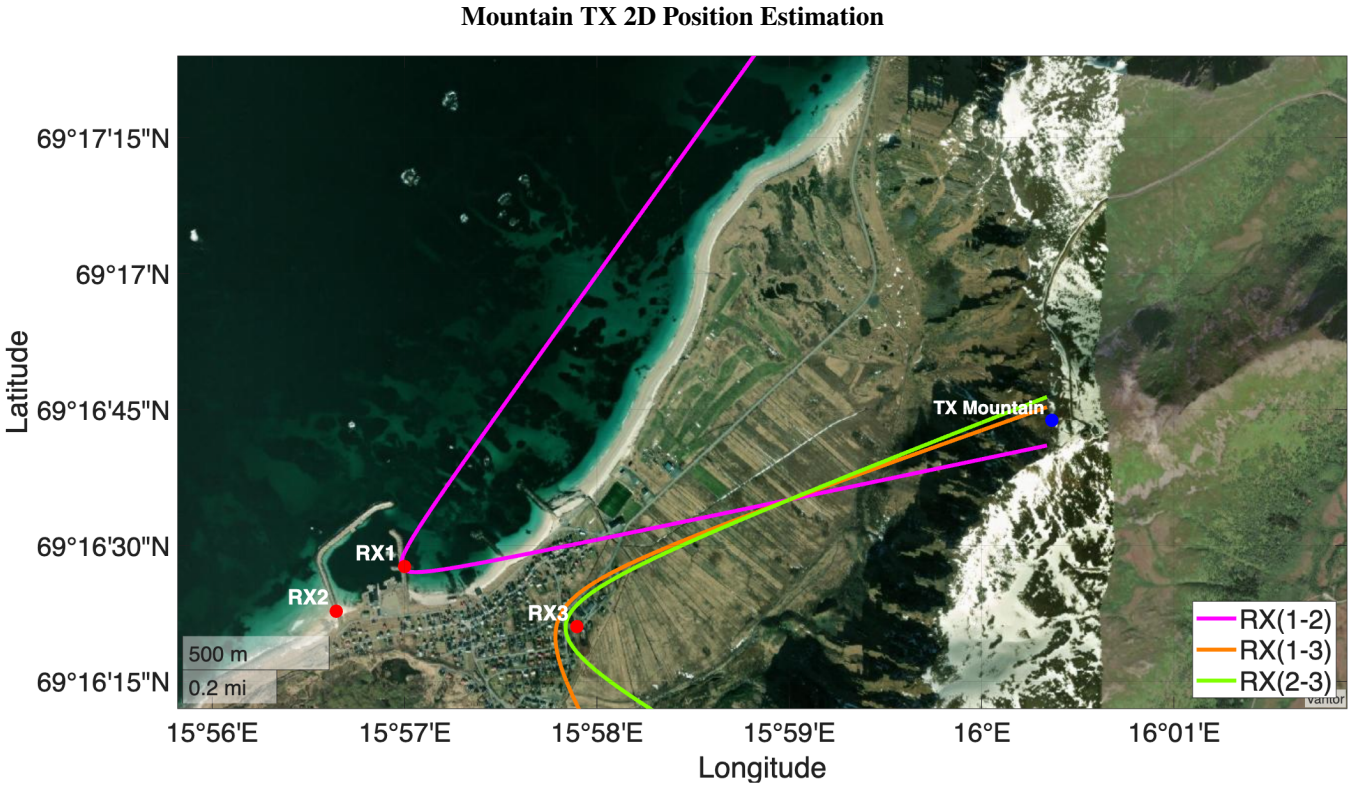


**Figure 7:** Mean hyperbola solution using 2nd order.

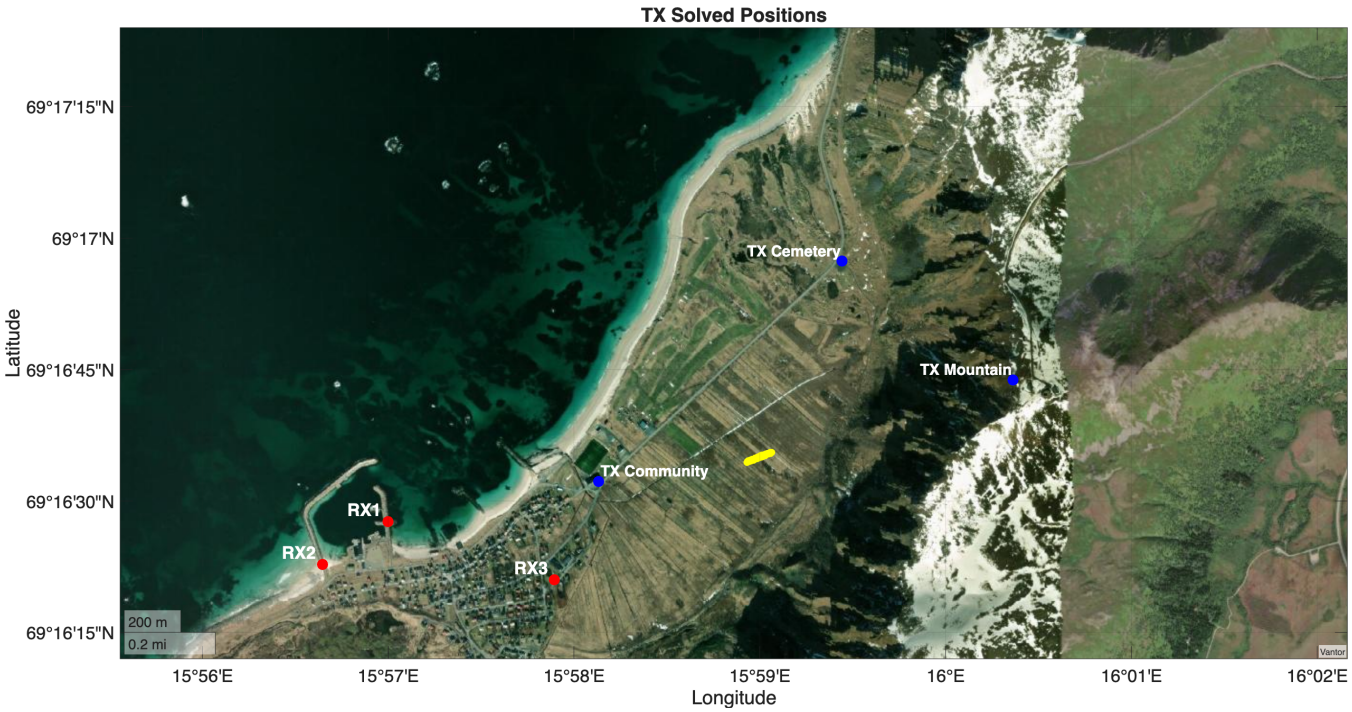


**Figure 8:** NLLS solution at each epoch (yellow) using 2nd order.

## 2. Community Transmitter (Spoofer)

The Community transmitter operated as a spoofer in close proximity to the receiver network, with $C/N_0$ values shown in Figs. 9 - 12 and an example of the intended spoofed path for one captured receiver in Fig. 13. The surrounding receiver geometry provided favorable HDOP and well-conditioned TDOA geometry, and the NLLS solver converged to a tight cluster of estimates within a few meters of the known transmitter location (Figs. 14 – 15), providing a controlled end-to-end validation of the methodology under near-ideal geometric conditions.

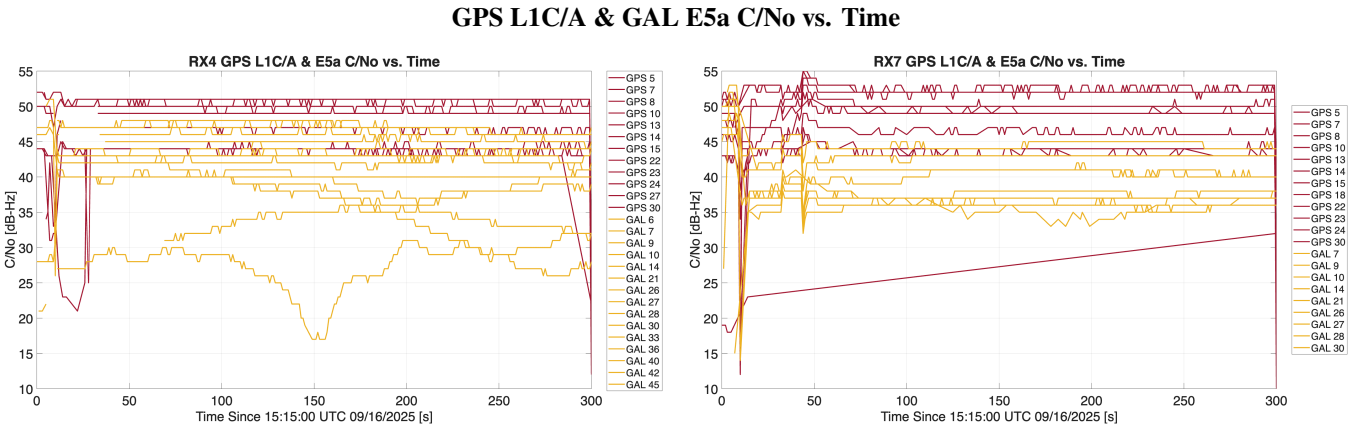


**Figure 9:** RX4 C/No during spoofing from the community.

**Figure 10:** RX7 C/No during spoofing from the community.

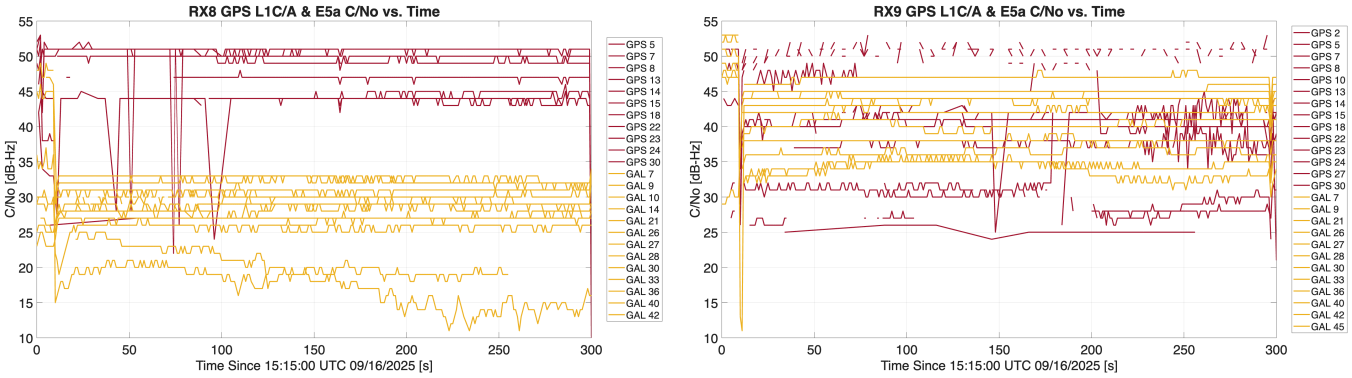


**Figure 11:** RX8 C/No during spoofing from the community.

**Figure 12:** RX9 C/No during spoofing from the community.

## 3. Cemetery Transmitter (Spoofer)

The Cemetery transmitter, like the Mountain transmitter, was located outside the receiver network. The spoofed path is seen in Fig. 19 for a receiver that was captured during this period. The $C/N_0$ measurements are sufficiently consistent to support the classification framework as seen in Figs. 16 – 18, and the individual TDOA position estimates shown in yellow are directionally coherent with the known transmitter location as seen in Figs. 20 and 21. However, the estimates exhibit a large spatial spread, which is attributable to poor horizontal dilution of precision (HDOP) arising from the combination of the transmitter's external position and the receiver network geometry. When the transmitter lies outside the network, the hyperbolic lines of position become nearly parallel in the vicinity of the true solution, amplifying the sensitivity of the NLLS solver to measurement noise. As with the Mountain scenario, a reliable point solution cannot be recovered, but the aggregate direction of arrival is consistent across estimates, reinforcing the utility of a direction-of-arrival interpretation in geometrically constrained configurations.

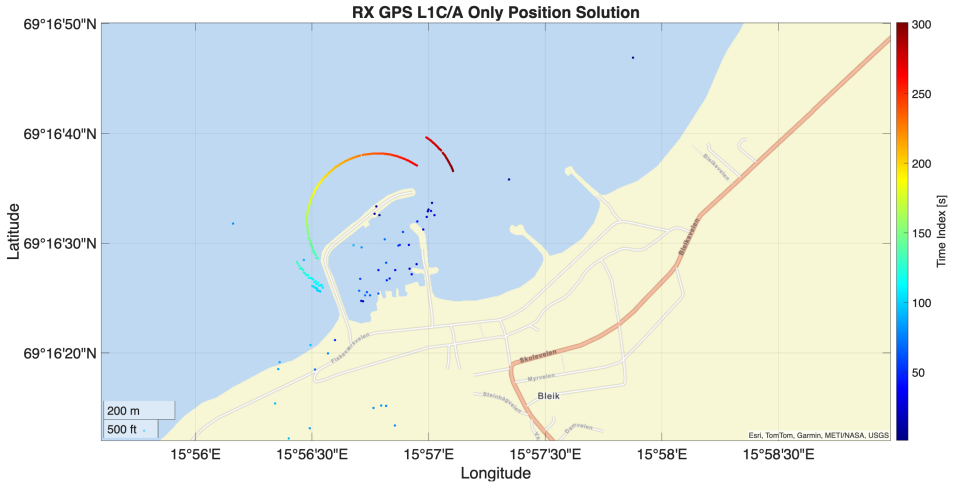


**Figure 13:** Example of full RX capture from community with GPS L1C/A measurements only as a function of time (simulated circle route).

**GPS L1C/A & GAL E5a C/No vs. Time**

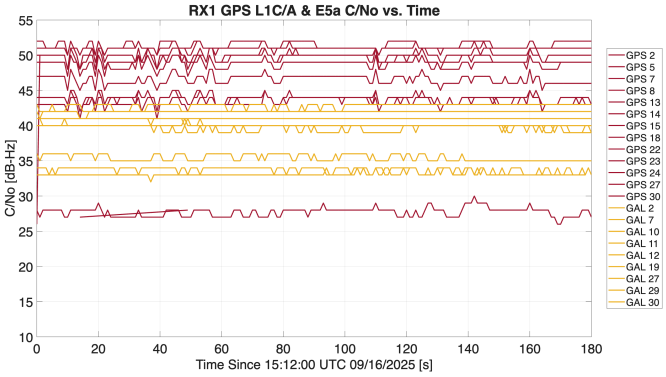


**Figure 16:** RX1 C/No during spoofing from the cemetery.

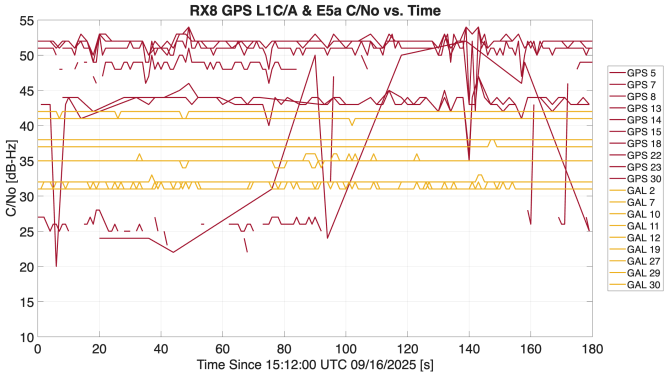


**Figure 17:** RX8 C/No during spoofing from the cemetery.

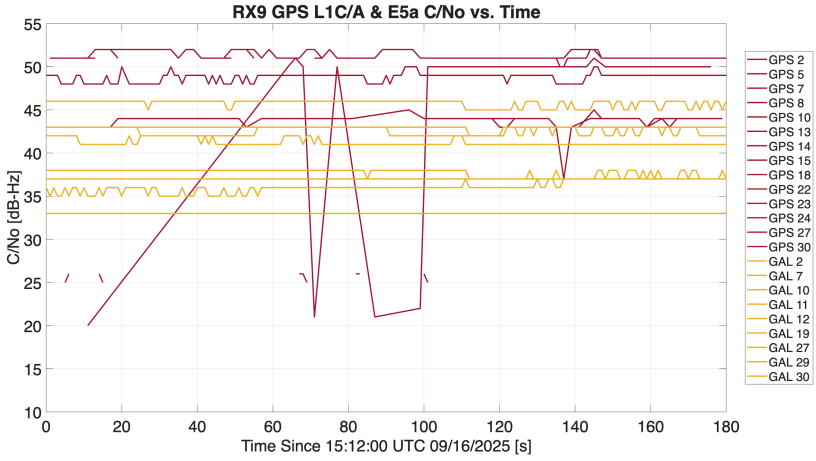


**Figure 18:** RX9 C/No during spoofing from the cemetery.

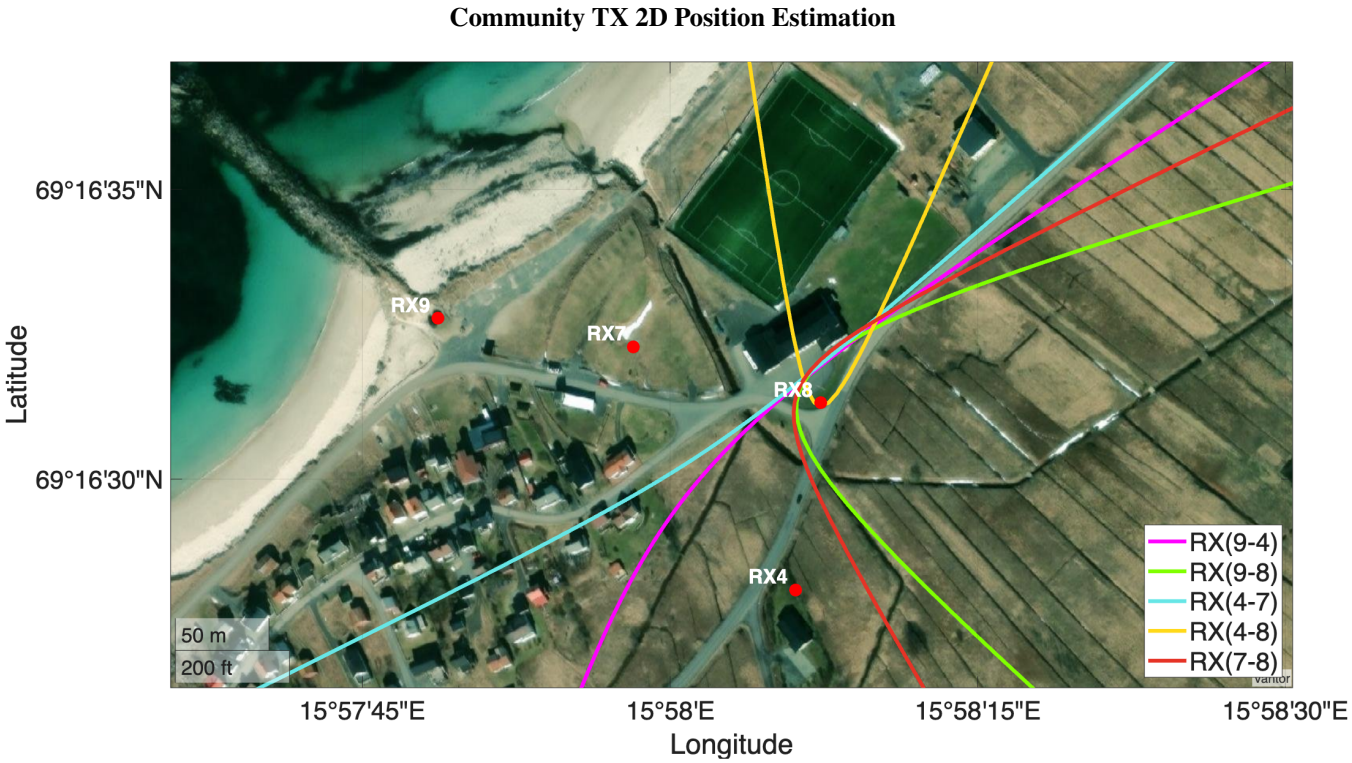


**Figure 14:** Mean hyperbola solution using 2nd order.

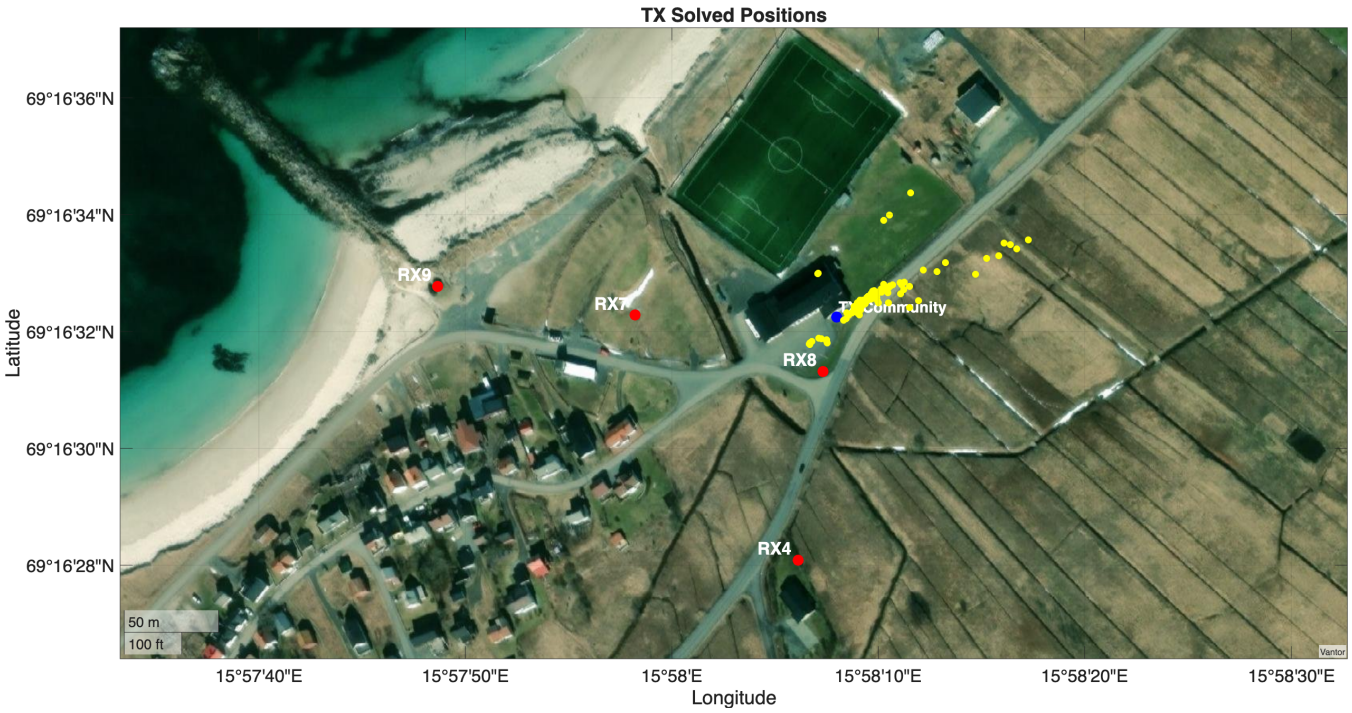


**Figure 15:** NLLS solution at each epoch (yellow) using 2nd order.

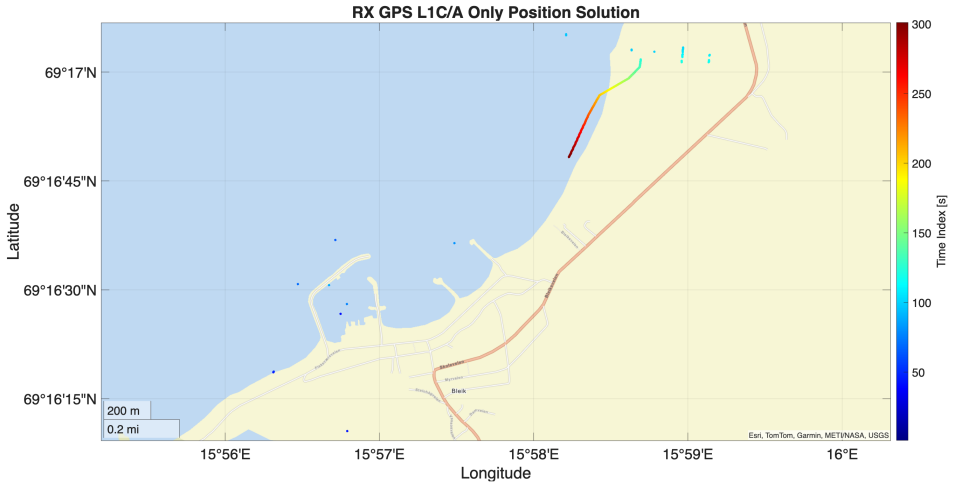


**Figure 19:** Example of full RX capture from cemetery with GPS L1C/A measurements only as a function of time (simulated driving route).

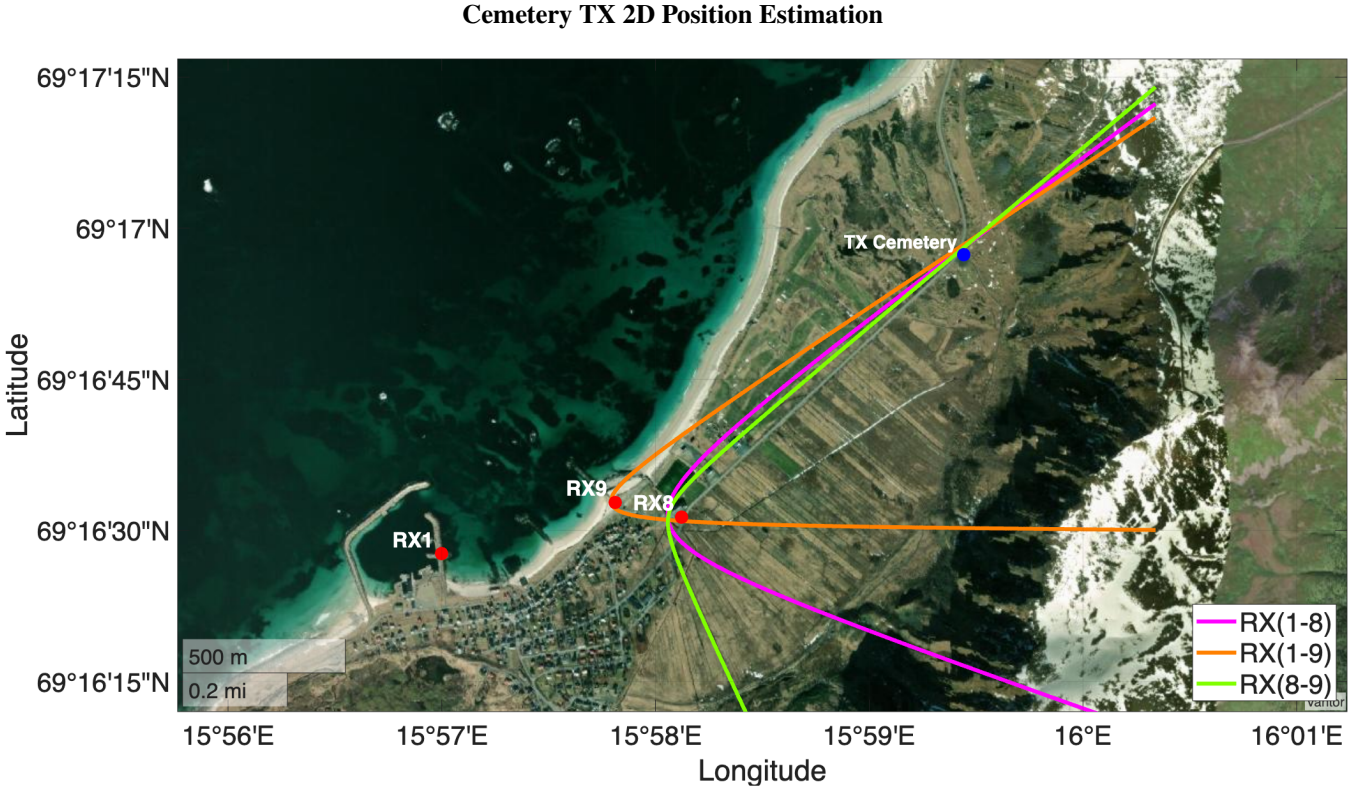


**Figure 20:** Mean hyperbola solution using 2nd order.

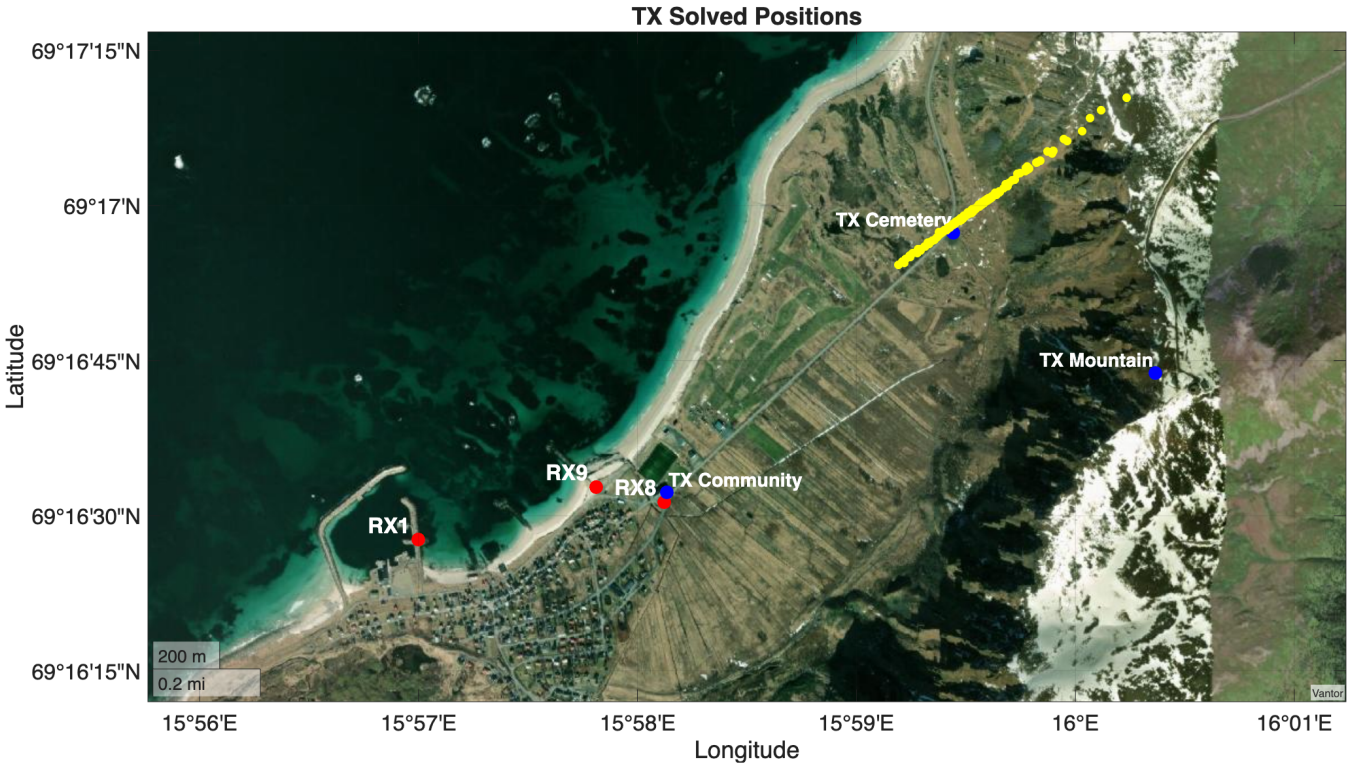


**Figure 21:** NLLS solution at each epoch (yellow) using 2nd order.

**Table 1:** Mountain TX 2D Error

| Method | Mean | Median | Std |
|---|---|---|---|
| 2nd order | 577.05 | 576.14 | 19.52 |
| 3rd order | 577.07 | 576.14 | 19.56 |
| Parametric | 579.99 | 580.10 | 19.58 |

**Table 2:** Community TX 2D Error

| Method | Mean | Median | Std |
|---|---|---|---|
| 2nd order | 23.89 | 18.38 | 19.03 |
| 3rd order | 21.36 | 17.70 | 15.06 |
| Parametric | 24.21 | 18.06 | 17.84 |

**Table 3:** Cemetery TX 2D Error

| Method | Mean | Median | Std |
|---|---|---|---|
| 2nd order | 134.32 | 103.60 | 115.27 |
| 3rd order | 134.23 | 103.84 | 114.51 |
| Parametric | 132.42 | 105.57 | 99.37 |

### 4. Future Directions for Multi-Constellation Integration

Tables 1 – 3 summarize the localization results obtained using each of the three $\sigma_\rho(C/N_0)$ models. The results are broadly consistent across all three models, which is expected given that the majority of observed signals maintained moderate-to-high $C/N_0$ values throughout the test periods. In this regime, the second-order polynomial, third-order polynomial, and parametric fits produce similar noise variance estimates, and meaningful differences between the models are only expected to emerge when a significant proportion of measurements fall in the low $C/N_0$ range where the curves diverge and extrapolation becomes necessary.

The present analysis relied exclusively on GPS L1 C/A and Galileo E5a measurements. Incorporating additional constellations and frequencies would increase the number of available measurements, improve the probability of forming valid quartets, and provide redundancy in both the classification and localization steps. A natural extension of this work would be to compute an independent TDOA position solution for each constellation-frequency pair and compare each solution against the known receiver positions. To assess internal consistency, a $\chi^2$ goodness-of-fit test could be applied to each solution residual vector. It should be noted, however, that a fully spoofed position solution would also pass such a test, since a self-consistent but incorrect solution is indistinguishable from a correct one under a $\chi^2$ criterion alone. This motivates the cross-constellation comparison as the primary discriminator: a spoofed signal set and an authentic signal set processed independently should produce inconsistent position solutions, whereas two authentic sets should agree. This approach would substantially reduce the combinatorial search space of the quartet-based framework, replacing the exhaustive receiver-satellite enumeration with a more tractable constellation-level hypothesis test.

## VI. CONCLUSION

This work presented a pseudorange double-difference framework for discriminating between authentic and spoofed GNSS signals across a network of spatially separated receivers, and demonstrated its application to spoofer transmitter localization via TDOA. A $C/N_0$-based noise model was empirically derived from rooftop measurements at Stanford University, and a geometric separation condition $|L| \geq 4\sigma_x$ was enforced to ensure reliable hypothesis discrimination. The methodology was validated against three spoofing scenarios from Jammertest 2025 using GPS L1 C/A and Galileo E5a measurements. For the Community transmitter, the NLLS solver converged to within a few meters of the known transmitter location. For the Mountain and Cemetery transmitters, which lay outside the receiver network, point solutions could not be recovered due to poor HDOP, though the TDOA estimates were directionally consistent with the known transmitter locations in both cases. Key limitations include the zero-mean Gaussian noise assumption, the heuristic choice of decision threshold, the restriction to two of sixteen possible quartet hypotheses, and the requirement for at least two authentic and two spoofed signals to be present simultaneously.

## VII. FUTURE WORK

A primary direction for future work is the development of a network-wide signal classification framework. The current brute-force enumeration of receiver and satellite combinations is computationally redundant, may reclassify the same signal multiple times, and leaves a subset of measurements inconclusive due to poor quartet geometry or insufficient signal strength. The goal is to replace this with a unified hypothesis test over all measurements across all receivers simultaneously, jointly estimating which signals are authentic and which are spoofed while concurrently estimating the transmitter location. The quartet-based framework developed here provides a natural foundation for this: quartets that yield high-confidence classifications can be used to anchor the network-wide solution, constraining the hypothesis space and reducing the number of unknowns that must be resolved for ambiguous measurements. This tighter coupling between classification and localization (rather than treating them as sequential steps) is expected to improve both the completeness of the signal labeling and the accuracy of the transmitter position estimate,

particularly in geometrically challenging configurations where the transmitter lies outside the receiver network.

Incorporating additional GNSS constellations and frequencies will further improve measurement redundancy and enable a complementary constellation-level $\chi^2$ consistency test, which would substantially reduce the combinatorial search space and provide an independent check on the classification output. Finally, the systematic positional bias observed in the Mountain and Cemetery scenarios will be investigated with the aim of recovering reliable point solutions for transmitters located outside the receiver network.

## ACKNOWLEDGEMENTS

We want to acknowledge and thank CARNATIONS, The Aerospace Corporation, and the FAA for sponsoring this research.

## REFERENCES

Babcock-Chi, Jade, Lo, Sherman, Chen, Yu-Hsuan, Blanch, Juan, Walter, Todd, "TDOA-Based Spoofing Source Localization Using Multi-Constellation Unsynchronized Receivers,"Proceedings of the 38th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2025), Baltimore, Maryland, September 2025, pp. 778-796.https://doi.org/10.33012/2025.20369

Hofmann-Wellenhof, B., Lichtenegger, H., Wasle, E. (2008). GNSS – Global Navigation Satellite Systems: GPS, GLONASS, Galileo more. Springer.

Humphreys, T. E. (2013). Detection strategy for cryptographic GNSS anti-spoofing. IEEE Transactions on Aerospace and Electronic Systems, 49(2), 1073–1090. https://doi.org/10.1109/TAES.2013.6494372

A. J. Jahromi, A. Broumandan and G. Lachapelle, "Gnss signal authenticity verification using carrier phase measurements with multiple receivers," 2016 8th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC), Noordwijk, Netherlands, 2016, pp. 1-11, doi: 10.1109/NAVITEC.2016.7849323

Jafarnia-Jahromi, A., Broumandan, A., Nielsen, J., Lachapelle, G. (2012). GPS vulnerability to spoofing threats and a review of antispoofing techniques. International Journal of Navigation and Observation, 2012, Article 127072. https://doi.org/10.1155/2012/127072

Medina, D., Gibson, K., Ziebold, R., Closas, P. (2018). Determination of pseudorange error models and multipath characterization under signal-degraded scenarios. In Proceedings of ION GNSS+ 2018

Misra, P., Enge, P. (2012). Global Positioning System: Signals, measurements, and performance (2nd ed.). Ganga-Jamuna Press.

Motella, B., Pini, M., Cavaleri, A., Fortuny-Guasch, J., Falletti, E., Lo Presti, L. (2011). Performance assessment of spoofing countermeasures based on signal quality monitoring for civil aviation. In Proceedings of the 24th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2011) (pp. 2609–2619).

Psiaki, M. L., Humphreys, T. E. (2016). GNSS spoofing and detection. Proceedings of the IEEE, 104(6), 1258–1270. https://doi.org/10.1109/JPROC.2016.2526658

Qin, H., Gong, H., Ma, Y., Zhen, D. (2018). Spoofing detection algorithm based on pseudorange double differences for multi-antenna GNSS receivers. Sensors, 18(1), 247. https://doi.org/10.3390/s18010247

Tippenhauer, N. O., Pöpper, C., Rasmussen, K. B., Čapkun, S. (2011). On the requirements for successful GPS spoofing attacks. In Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS '11) (pp. 75–86). ACM. https://doi.org/10.1145/2046707.2046719

Wang, L., Groves, P. D., Ziebart, M. (2017). GNSS NLOS and multipath error mitigation using advanced multi-constellation consistency checking with 3D mapping. GPS Solutions, 21(2), 381–395. https://doi.org/10.1007/s10291-016-0525-3

Zhu, N., Yang, Y., Li, J. (2019). GNSS spoofing network monitoring based on differential pseudorange. Sensors, 19(21), 4769. https://doi.org/10.3390/s19214769