MAPT: A Near Real-Time Adaptive Multipath Modeling and Clustering Algorithm for GNSS Urban Positioning

Julian Gutierrez
Safety Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia, United States
julian.gutierrez@nasa.gov

Russell Gilabert
Safety Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia, United States
russell.gilabert@nasa.gov

J Tanner Slagel
Safety Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia, United States
j.tanner.slagel@nasa.gov

Evan Dill
Safety Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, Virginia, United States
evan.t.dill@nasa.gov

Pau Closas
Associate Professor
Northeastern University
Boston, Massachusetts, United States
closas@ece.neu.edu

***Abstract*—Position estimation in urban canyons using Global Navigation Satellite Systems (GNSS) often suffers from reduced accuracy due to signal interactions with buildings. One example is multipath, where signals reflect off surfaces and introduce delays in the received signal (often resulting in pseudorange estimation errors). In this work, we introduce MAPT (Multipath Adaptive Positioning Techniques) to improve positioning accuracy. MAPT accomplishes this goal by integrating: 1) the NavQ line-of-sight (LOS) performance monitor, 2) a simplified multipath model within NavQ with a non-line-of-sight (NLOS) masking technique, 3) a satellite combination and position solution processor, and 4) a clustering algorithm. These techniques are collectively leveraged to identify trustworthy satellite signals for accurate positioning. MAPT evaluated across 5,475 data points from a representative urban canyon in Columbus, Ohio reduced the 95^{th} percentile of horizontal error from 47.8 meters (without processing) to 4.4 meters. Additionally, this was accomplished by processing each epoch in under 0.24 seconds on average.**

I. Introduction

The need for high-accuracy positioning is growing, particularly for future Unmanned Aerial Systems (UAS) and Urban Air Mobility (UAM) operations [1]. Global Navigation Satellite Systems (GNSS) continue to be essential for optimizing real-time traffic management and location-based services. However, urban environments present a unique set of challenges which can cause significant errors in positioning estimation, affecting airborne vehicles as well as ground users.

To improve GNSS integrity, systems like Receiver Autonomous Integrity Monitoring (RAIM) detect and mitigate GPS errors [2]. However, RAIM's effectiveness decreases in urban areas due in part to multipath effects [3].

In terms of accuracy and availability, integrating GNSS with other sensors can enhance accuracy; however, these methods are usually non-trivial and can still rely too heavily on GNSS techniques, highlighting the need for innovative solutions such as selective satellite pruning to enable UAM operations.

We build upon prior work introduced in [4] to improve the detection and mitigation of multipath effects in GNSS position estimation. MAPT (Multipath Adaptive Positioning Techniques) combines a NASA-developed GNSS quality prognostic tool called NavQ [5], [6], a multipath modeling approach, and a clustering algorithm to evaluate multiple satellite combinations and identify problematic measurements from satellites, especially those affected by multipath. This approach is optimized to evaluate a smaller subset of satellites while still retaining similar levels of accuracy observed in prior work.

MAPT helps advance four research objectives:

1) Improving position estimation accuracy in urban environments.
2) Leveraging an efficient multipath model that uses non-line-of-sight (NLOS) masking to weight the likelihood of multipath effects on line-of-sight (LOS) signals.
3) Utilizing machine learning for multipath mitigation that uses a clustering algorithm to cluster different positioning solutions and determine sets of satellites that are likely to be multipath-free, in a way that generalizes detection and exclusion algorithms in integrity monitoring to multiple faulty satellites.
4) Optimizing MAPT algorithm performance for real-time applications.

II. MAPT Algorithm Design

MAPT is a collection of techniques that integrate Machine Learning (ML) with traditional navigation algorithms. By leveraging the strengths of both methods, MAPT can adaptively mitigate the effects of multipath and provide more accurate location estimates.

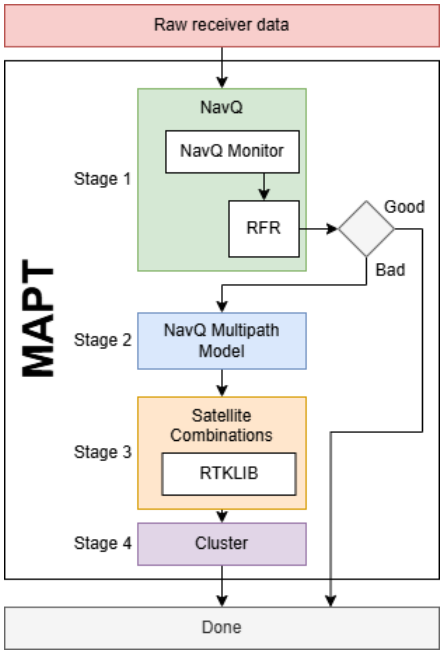

Fig. 1. MAPT Algorithm Flowchart.

As shown in figure 1, MAPT consists of four stages:

1) **NavQ Execution**: This stage leverages the NavQ prognostic service to determine which satellites are in LOS and collects multiple performance metrics. These performance metrics are fed into a statistical model to estimate the accuracy of the positioning solution obtained using the selected LOS satellites.
2) **NavQ Multipath Model Prediction**: In this stage, a simplified model of the physical features in the surroundings is used to predict the multipath effects for the GNSS signals. Additionally, the LOS multipath estimation is augmented with NLOS satellite multipath masks.
3) **Satellite Combinations and Positions Computation**: This stage adds the predicted multipath affected satellites to a list of satellites and produces combinations of them. Each combination is then processed by RTKLIB's solver [7] and a position solution generated.
4) **Cluster Determination**: In this final stage, the algorithm determines the most likely cluster of position solutions generated with suspected "multipath-free" satellites and computes a final position estimate using the centroid of the cluster.

The MAPT algorithm is designed to use RTKLIB's single point positioning (SPP) solver, which estimates the location of a single receiver by using the pseudorange measurements of different GNSS satellites. Unlike some other navigation algorithms that rely on previously computed solutions, MAPT does not reuse previous results; instead, each solution is independently computed. Using these techniques allows for the treatment of each epoch in isolation to determine the effectiveness of the algorithm.

The output from the statistical model in Stage 1, a confidence metric, is used to evaluate the position solution using only LOS satellites. If the confidence associated with NavQ's solution exceeds a user-defined threshold, then that solution is used and the rest of the algorithm is skipped, saving computational time. A detail description of each stage is provided in the following sections.

A. Stage 1 - NavQ Execution

The NavQ execution stage utilizes the NavQ GNSS performance monitor to distinguish between LOS and NLOS satellites for each epoch [5], [6]. Identified NLOS satellites are subsequently removed from the list of satellite signals that will be analyzed in later stages.

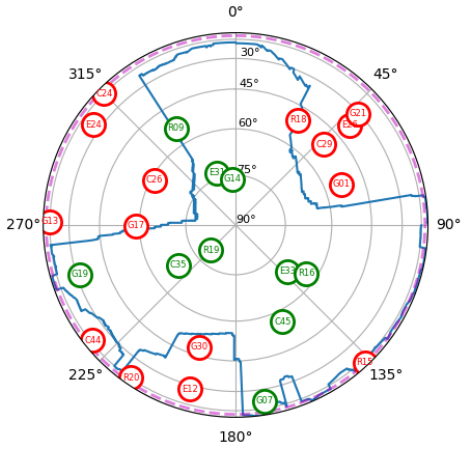

(a) Skyplot

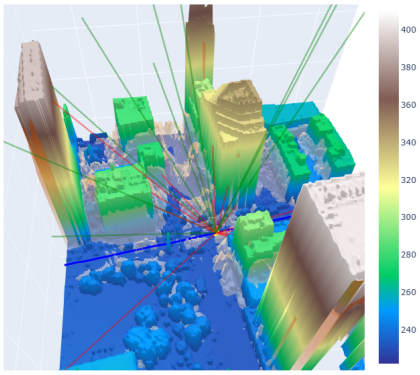

(b) 3D visualization

Fig. 2. Example output from NavQ [6]. (a) shows the skyplot indicating the blockages caused by the buildings surrounding the receivers position as well as the positions of the satellites in the sky and whether they are blocked (red) or there is direct LOS to them (green). (b) is a 3D visualization for the same location and displaying the rays in the direction of the satellites using the same color scheme. North is directed to the right of the figure.

As discussed in the literature, LOS filtering has been proven to be an effective method for reducing errors caused by NLOS multipath interference [5]. To perform LOS/NLOS filtering,

NavQ utilizes a combination of the fused GNSS/inertial solution, a 3D model of the urban area (i.e., physical structures such as buildings), and satellite geometry information. The fused GNSS/inertial solution is used to determine where the GNSS satellites are in the sky relative to the local building geometry [8]. This information is then used to predict LOS/N-LOS status of each satellite to identify any potential NLOS multipath sources, as shown in Figure 2. By removing these interfering signals from the navigation solution, the quality of the position estimate can be significantly improved.

To evaluate the quality of NavQ's solution, we developed a confidence metric based on ML techniques. An ML solution was chosen because the quality of NavQ's solution depends on several factors in a discontinuous and non-linear way that is not easy to model. A ML solution provides a way to model this complex behavior and supply a level of confidence that is computationally feasible at run-time.

Specifically, a random forest regressor (RFR) model was trained to predict the quality of the proposed NavQ solution which we defined as $1/(1 + error)$, where $error$ is the Euclidean distance between the estimated and ground truth positions (the method to establish the ground truth is described in Section III). Modeling the error this way bounds the quality term to values between 0 and 1. During training, the regressor model leverages the following three key features related to the LOS satellites to calculate the solution confidence (where higher confidence indicates a higher quality, more reliable solution):

- **Number of LOS satellites**
- **Percentage of sky visibility**: The fraction of the sky that is visible to the receiver, which can suggest potential multipath effects
- **Horizontal Dilution of Precision (HDOP) of LOS satellites**

Solution confidence uses these three metrics based on an assumption that the combination of more satellites in view, higher sky visibility, and low HDOP is most likely to produce a trustworthy position solution. Conversely, if visibility is low or HDOP is high for LOS satellites, these conditions suggest a higher likelihood of multipath effects and therefore a situation in which the full MAPT solution is more likely to be better.

To optimize the performance of the regressor model, we employed a hyper-parameter grid search using the GridSearchCV function from scikit-learn library [9]. This approach was selected because it avoids overfitting the parameter space by evaluating the models using five-fold cross-validation [10]. Additionally, custom scoring is used to evaluate confidence prediction and penalize overconfidence (i.e., high accuracy in predicted confidence when the error is small).

B. Stage 2 - NavQ Multipath Model Prediction

This stage enhances NavQ by assigning a likelihood score to each satellite and quantifying the probability of signal reflections coming from all azimuth directions. This parameterized model takes into account critical factors that influence multipath effects, including:

- **The angle of the signal bounce relative to the reflective surface**: A steeper angle typically results in weaker multipath signals.
- **The distance of the bounce relative to the receiver**: Closer bounces tend to produce stronger multipath signals.
- **The satellite's azimuth and elevation angles**: These factors determine the directionality of the signal and its potential interaction with surrounding surfaces.

By incorporating these parameters, NavQ's multipath model provides a higher fidelity assessment of the likelihood and severity of multipath effects for each satellite. This approach is deliberately simplified, focusing on a basic heuristic that rapidly identifies potential satellite measurements with significant multipath interference. This trade-off between simplicity and accuracy enables the algorithm to provide a fast and preliminary assessment of potential issues.

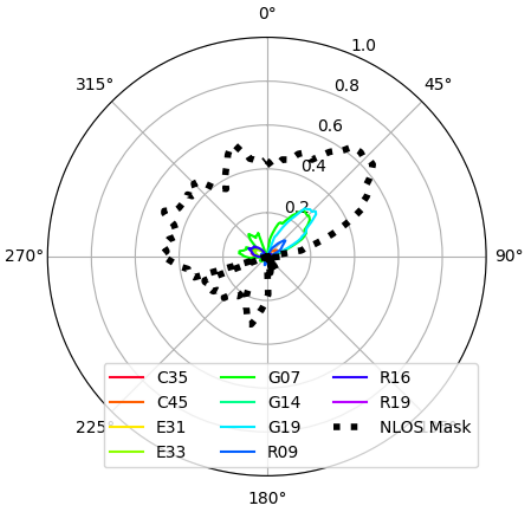

Fig. 3. Example output from NavQ's multipath model prediction for the LOS satellites plotted as a skyplot. The radius indicates the likelihood of multipath coming from a given azimuth direction and the color indicates the values for each satellite. The NLOS mask is the mask generated by using the maximum likelihood from all of the NLOS satellites the receiver observed.

Building on the NavQ multipath model's predictive capability, a mask of NLOS satellite likelihoods is created, retaining the maximum value among NLOS satellites observed by the receiver. This mask is then multiplied with the LOS satellite likelihoods to weigh potential multipath directions and likelihoods. Signal and visibility information from observed NLOS satellites is leveraged to influence the probability and direction of multipath effects on LOS satellites, which are subject to similar environmental factors as NLOS signals and may thus exhibit comparable interference patterns.

At the end of this stage, NavQ has compiled a list of satellites for which it estimates the highest likelihood of multipath interference. The number of satellites on this list depends on (1) the total number of satellites, (2) the NLOS mask, and (3) the magnitude of the multipath interference likelihood estimate (MILE).

Figure 3 shows an example of the output NavQ generates for the multipath model. In this example, the model suggested a higher probability of multipath for satellites G19 and G07, as demonstrated by the larger radius of their respective curves in the figure. After further processing of this epoch, we observed and confirmed a large multipath error on G19.

C. Stage 3 - Satellite Combinations and Positions Computation

A satellite combination refers to a set of satellites whose transmissions may be used to compute a position estimate. This stage aims to determine which combinations of satellites (and their corresponding position solutions) should be passed on to the clustering algorithm for further processing. By leveraging NavQ's predictive capability to identify LOS satellites, the state space of possible solutions can be reduced and the efficiency of the combination process can be improved. Additionally, NavQ's multipath model identifies the least reliable LOS satellites to use in these combinations.

The stage begins by using the parameter, N, which defines the number of satellites available from which to generate combinations. The number of combinations grows exponentially with the value of N, requiring careful selection of this parameter to prevent prolonged processing times. When selecting specific satellites to be used to generate combinations, satellites with the highest multipath likelihood from the previous stage are prioritized first. Any remaining spots are filled with the lowest elevation satellites, which can have a substantial effect on the final position solution if affected by multipath [11]. Any remaining LOS satellites (e.g. the high elevation angle satellites) remain fixed and are appended to all the combinations. These satellites are referred to here as keystone satellites. This approach creates a diverse set of satellite combinations without creating too many satellite combinations and requiring more computational time.

We utilize RTKLIB[1] to obtain position solutions for all combinations of satellites. Some combinations may not generate valid position solutions due to insufficient satellite count or solver failure. In rare cases, the algorithm can incorporate high elevation NLOS satellites and retest the new set of combinations to create sufficient valid position solutions. However, if the algorithm incorporates enough NLOS satellites and it is still unable to find a valid position solution, the execution is terminated and a result of "no successful position solution found" is reported.

D. Stage 4 - Cluster Determination

This stage - the clustering algorithm - determines the combination of satellites that yield a more accurate position solution. The algorithm produces multiple solutions based on these combinations and identifies the densest collection of position solutions, where the dense regions are most likely to exclude reflected satellite signals.

[1]We incorporated an epoch aggregation feature into the demo5 version of RTKLIB [7], allowing for simultaneous processing of satellite combinations and lower execution time.

We selected DBSCAN as the clustering algorithm technique used in this stage. The DBSCAN algorithm is a density-based clustering technique that groups data points based on minimum group size and relative proximity, denoted as epsilon (ϵ) [12]. DBSCAN is adapted for use in this stage by defining a target number of position solutions for each cluster. To achieve this, ϵ is adjusted until a cluster with sufficient size is found.

The process begins by estimating the number of potential combinations of satellites, that is, the number of points in a cluster, that will yield valid position solutions. These estimates are based on: 1) the total number of satellites being processed; 2) the number of satellites believed to have an error (defaulted to half the number of satellites being analyzed, rounded down); 3) the minimum size of a subset (set to 5 satellites); and 4) the number of keystone satellites that will appear in all combinations. The minimum cluster size is determined by first finding the ratio between the estimated potential combinations and the number of valid combinations if 2) was set to zero (no satellites have multipath). This ration is multiplied by the total number of valid combinations found to obtain the final minimum cluster size.

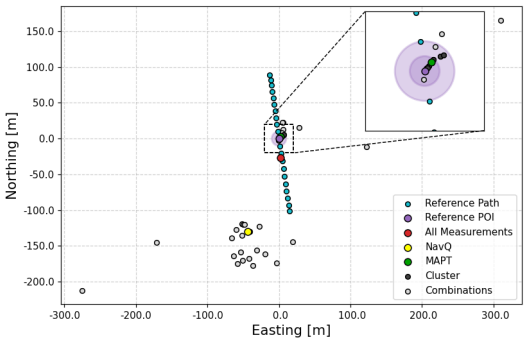

Fig. 4. Position solutions generated from: 1) using all available signals to the receiver (red), 2) using only the LOS satellite signals determined by NavQ (yellow), 3) the final solution determined by MAPT (green), 4) the cluster selected by MAPT (black), and 5) the different combinations of satellites MAPT tested (white). The reference POI is marked with purple circles with a 5 and 10 m radius, and the reference path of the vehicle before and after the selected epoch is marked in teal.

This clustering approach enables the systematic identification of the densest cluster, which is assumed to contain position solutions from satellite combinations that have not been affected by multipath interference. Building on prior work wherein non-multipath position solutions were observed to be generally well-grouped [6], any outlier position solution *not* found within this dense cluster is assumed to have a satellite combination containing one or more erroneous measurements due to multipath. Therefore, the final position solution and its corresponding satellites are ultimately determined by selecting the solution closest to the centroid of the densest cluster.

Figure 4 shows an example of the position of the different solutions MAPT can produce. In this example, the solution selected by MAPT is closer to the true position than NavQ's solution or the solution using all of the received signals (all measurements) as it is the closest to the true position (marked

with purple). This was a result of a dense cluster being found close to the true position, allowing MAPT to select a combination of satellites that correctly gets rid of the satellite affected by multipath, which in this case was G19.

III. Methodology

This study evaluates the performance of MAPT using data collected from both Monte Carlo simulations and real-life operations. Results provide insight into its limitations and effectiveness across a broad span of conditions.

The physical setup to collect the operational data consisted of a SparkFun u-blox F9R breakout GNSS receiver, which was chosen for its ability to track L1 and L2 signals from multiple satellite constellations (including GPS, Galileo, GLONASS, and Beidou) and generate a GNSS/inertial integrated solution. The receiver was connected to an ArduSimple simpleANT2B multiband GNSS antenna, which was deployed on the roof of a car using a magnetic base.

The receiver was connected to a laptop, which both powered the receiver and logged the GNSS data, including the raw pseudorange measurements, every second. The laptop also recorded additional metadata for each epoch such as timestamp, latitude, longitude, altitude, and velocity.

The downtown area of Columbus, Ohio (Figure 5) was the selected location for the data collection due to its moderately sized urban canyon, which is representative of a typical city environment. The vehicle was driven through the same route in the downtown area eight times over two days, with four test drives each day, collecting a total of 5475 epochs of data.

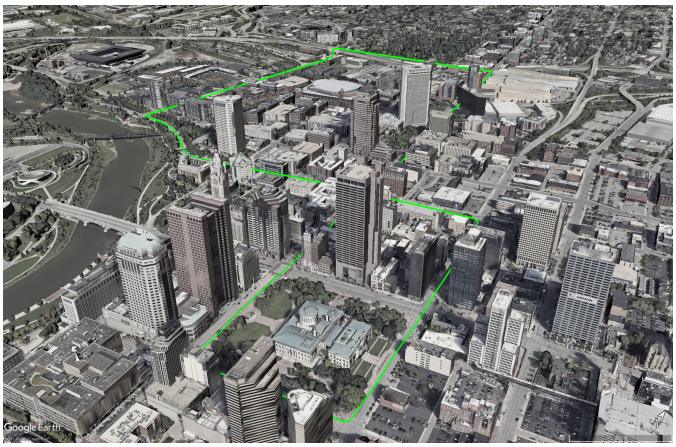

Fig. 5. Google Earth view of the city Columbus, Ohio with the reference ground track used for data collection (in green).

A ground-truth reference dataset was derived using the fused GNSS/inertial measurements generated by the F9R receiver, no postprocessing done. The position reference point of the antenna on the vehicle was established at an elevation of 2 m AGL (Above Ground Level), with mean sea level as the datum.

IV. Results

This section presents the results of the Monte Carlo analysis and real-world testing performed using MAPT.

A. MAPT Performance in Removing Satellites with an Erroneous Estimated Pseudorange

In order to assess the performance of MAPT, known errors were intentionally introduced into the pseudorange measurements of satellites. Throughout the remainder of the paper, sources of these modified pseudorange measurements will be referred to as satellites with an erroneous estimated pseudorange (SEEPRs). Next, a Monte Carlo simulation was conducted to evaluate the performance of the algorithm in detecting and correcting these SEEPRs.

The simulation utilized data from the NOAA Continuously Operating Reference Station (CORS) [13] NYBP, located in Battery Park, New York. This CORS station was programmed to process both GPS and GLONASS signals at 30-second intervals over a full day's duration for a total of 2,880 epochs. To simulate errors, the pseudorange of randomly selected satellites was modified using a Gaussian distribution with a mean error of 50m and variance of 10m. The selection of satellites to be modified was randomized with a probability of selection that was inversely proportional to each satellite's elevation angle. The RINEX file was altered to incorporate these errors and subsequently processed by MAPT. The horizontal error was computed at each epoch using the 2D Euclidean distance between the MAPT-derived position solution and the true CORS station location.

Simulation results consisted of a total of 144,000 epochs, providing an evaluation of the algorithm's performance under various conditions including: 1) varying number of LOS satellites, 2) varying the number of satellites where an artificial error was introduced in the estimated pseudorange, and 3) different geometric positions of the satellites in the sky. MAPT was configured to produce combinations for up to seven satellites, with any additional satellites being designated as the keystone satellites.

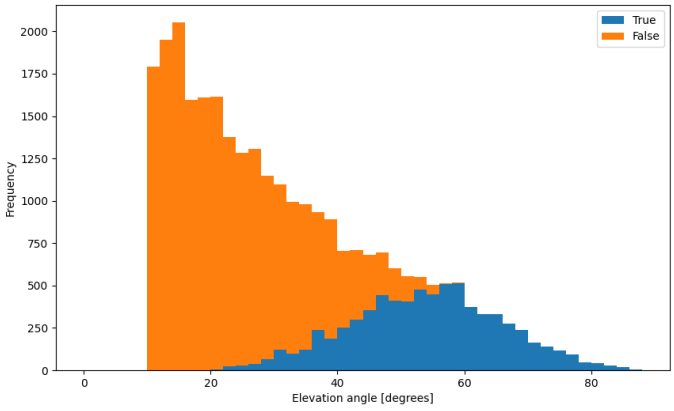

Fig. 6. Elevation angle histogram of the SEEPR for the subset of epochs where the error was introduced into only one satellite. The classification used in the legend indicates whether the affected satellite was a keystone satellite (blue) or not (orange).

Figure 6 presents the distribution of satellites where a pseudorange error was introduced based on their elevation angle. This figure reflects simulating multipath effects wherein

lower-elevation satellites are more susceptible to error and thus prioritized. The count shown in this figure represents the subset of Monte Carlo runs where a single-satellite-error was introduced. The colors used in the figure indicate the classification of satellites as either keystone satellites (blue), which MAPT is unable to correct, or satellites processed by MAPT (orange).

All simulated epochs were grouped by the number of SEEPRs. These groups are compared using an empirical cumulative distribution function (ECDF) of the horizontal error. Figure 7 displays the ECDF of MAPT's horizontal error as the number of SEEPRs increases. When no errors were introduced to the satellites (red curve), the ECDF curve exhibits the expected performance of pseudorange-based positioning, and all the horizontal errors are within 5 m.

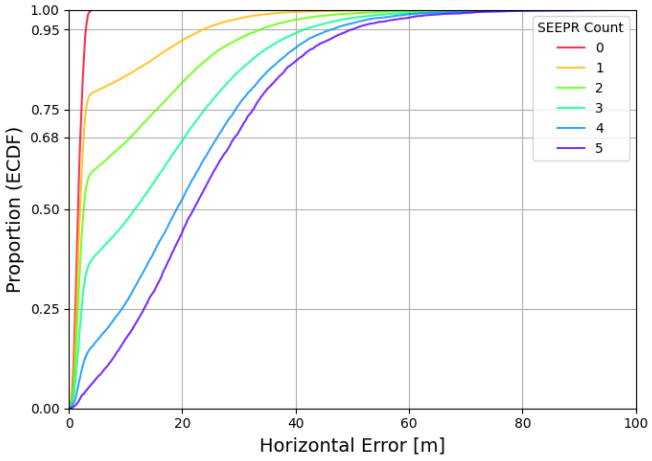

Fig. 7. Empirical cumulative distribution function (ECDF) of the Monte Carlo simulation results for horizontal error when varying the number of SEEPRs.

However, as the number of SEEPRs grows, a notable "elbow" appears in the ECDF curve. This elbow signifies the point at which MAPT is unable to correct these errors, resulting in a more gradual slope. The observed elbows in this ECDF are a direct result of the Monte Carlo simulation modifying keystone satellites that MAPT cannot correct. As more satellites are introduced with errors, there is an increased likelihood that one or more will be classified as keystone satellites which are beyond MAPT's processing capacity. This effect highlights the importance of MAPT selecting the proper satellites to analyze. NavQ's multipath model can aid the algorithm in real scenarios to carefully select better suited satellites for analysis, but in these artificial cases, NavQ cannot provide useful information.

To provide additional insight of MAPT's performance, the subset of data where no errors were introduced into keystone satellites is examined. This isolates the algorithm's capabilities to situations where it can analyze the affected satellites. Figure 8 presents the ECDF curves of the subset of data from Figure 7 where none of the affected satellites are classified as keystone satellites. The analysis shows that MAPT is effective in eliminating all artificial errors when one or two satellites have introduced errors, as indicated by the lack of "elbow" in the graph. However, algorithmic performance deteriorates as the number of SEEPRs increases; indeed, if five or more satellites are modified with an error, the algorithm becomes ineffective in finding a reasonable position solution, as evidenced by the curve's gradual slope.

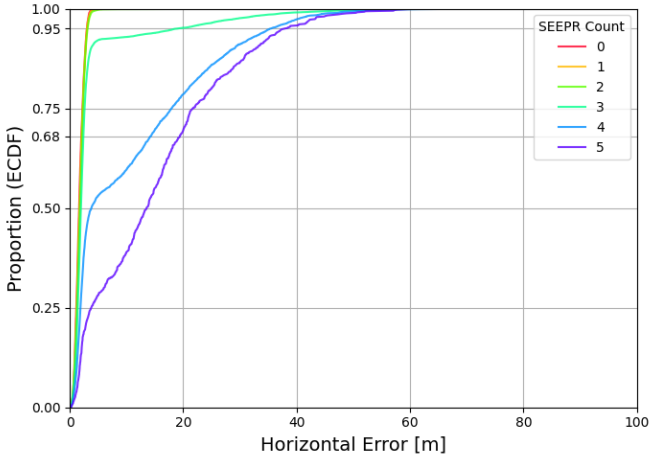

Fig. 8. Empirical cumulative distribution function (ECDF) of the Monte Carlo simulation results for horizontal error when varying the number of SEEPRs and no keystone satellites were affected.

Figure 8 shows that, when there are three SEEPRs, MAPT fails to find a position solution without using a SEEPR approximately 8% of the time. If MAPT analysis considers up to seven satellites then, despite the presence of three satellites where an error has been introduced in the estimated pseudorange, intuitively, MAPT should still be able to identify clusters with accurate position solutions among the remaining satellites. However, two phenomena were observed in epochs where the horizontal error was larger than 20 m (past the elbow point) which hint at why this might not be the case:

- When there are fewer LOS satellites, the number of valid position solutions decreases, making it more difficult to determine the impact of SEEPRs.
- As the SEEPR count increases, there is an increased likelihood that position solutions generated with one or more SEEPRs will converge into similar locations, creating a "erroneous cluster".

Figure 9 shows a breakdown of the three-SEEPRs curve from Figure 8, based on the total number of LOS satellites in the sky. As the total number of LOS satellites increases, the ECDF curve improves, and the proportion of epochs with low error also increases. To conclude, MAPT's performance improves significantly when analyzing seven satellites if there are at least 13 LOS satellites in the sky, allowing it to correct 95% of errors when three or fewer satellites contain multipath errors.

B. NavQ's Confidence Analysis

Figure 10 shows the RFR model output confidence as a function of the horizontal error in NavQ's solution (when

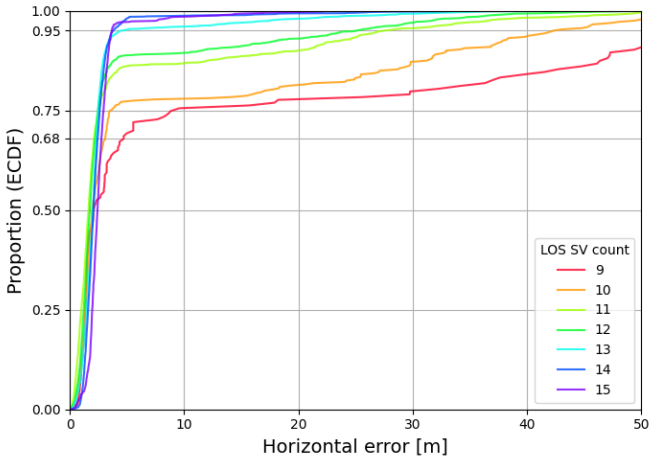

Fig. 9. Empirical cumulative distribution function (ECDF) of the Monte Carlo simulation results for horizontal error when three satellites have been modified and no keystone satellites were affected, plotted for the given number of LOS satellites in the sky.

only LOS satellites are used) based on the collected ground vehicle data described in Section III. The relationship between the confidence and horizontal error within NavQ's solution is evident, with higher confidence corresponding to lower errors while decreased confidence corresponds to increased errors.

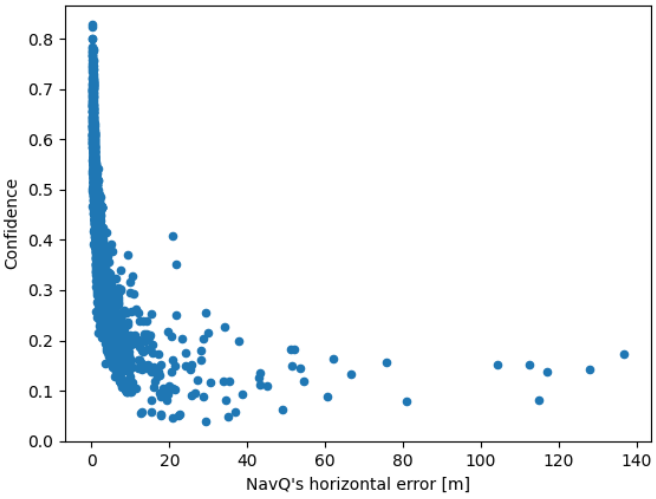

Fig. 10. MAPT confidence as a function of the horizontal error in the NavQ proposed solution when using only LOS satellites.

To determine the ideal confidence threshold for switching between the first-pass proposed NavQ solution and more resource-intensive MAPT solution, the threshold is varied from $0-1$ and compared the impact to the 95^{th} percentile of horizontal error across all datasets. The mean epoch processing time was used to evaluate the trade-off between accuracy and computational time. These results can be observed in Figure 11. Increasing the confidence threshold results in more cases where MAPT is fully executed, necessarily resulting in increased mean epoch processing time.

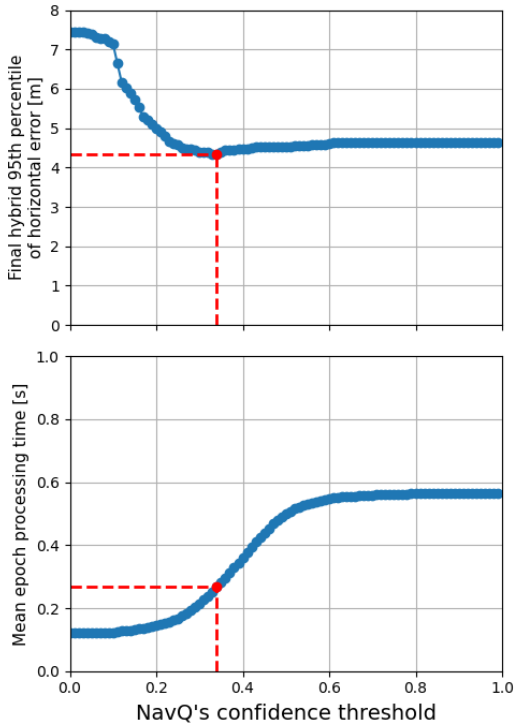

Fig. 11. 95th percentile of horizontal error and mean epoch processing time of MAPT as a function of NavQ's confidence threshold.

The results reveal that incorporating NavQ's position solution leads to marginally improved performance compared to MAPT solutions, particularly when using a NavQ confidence threshold of 0.33. In this configuration, a significant reduction in mean epoch processing time of 0.23 s was observed and a decrease in the 95^{th} percentile of horizontal error to 4.44 m. This marginal improvement compared to using MAPT at each measurement epoch indicates that NavQ outperforms MAPT in high-confidence scenarios. For the epochs with a confidence higher than 0.33, the horizontal accuracy using MAPT was found to be slightly lower than using NavQ alone in terms of the mean value (1.43 m and 1.37 m, respectively). In terms of variability, the difference was greater (standard deviation of 0.99 m and 0.69 m, respectively). As such, a NavQ confidence threshold of 0.33 was chosen for the analyses that follow.

C. Horizontal Error Analysis

Figure 12 shows the ECDF's for the collected real world datasets, graphing the 95^{th} percentile of the measured horizontal error for each of the three algorithmic approaches (i.e., when using all measurements vs. NavQ alone vs. MAPT). The use of all measurements resulted in a 95^{th} percentile error of 47.83 m. In contrast, employing NavQ reduced this error to 7.83 m, and MAPT further decreased it to 4.44 m, demonstrating a substantial reduction in the overall error versus using NavQ alone.

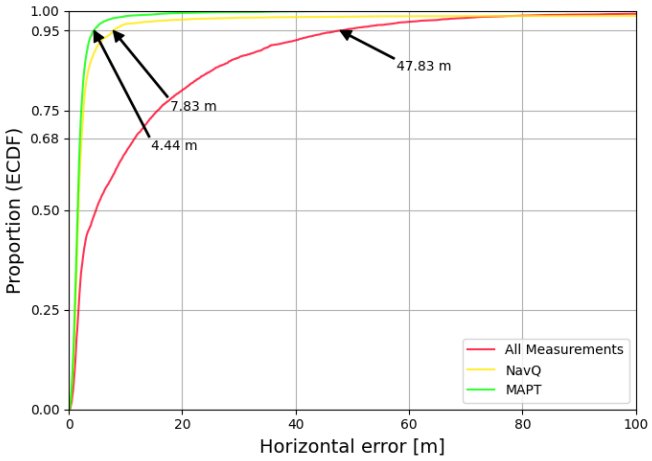

Fig. 12. Empirical cumulative distribution function (ECDF) of the horizontal error for different algorithms. (a) results for the full 8 datasets, which include 5445 epochs, and (b) results for the *urban canyon segment* of the path for the 8 datasets, which includes 1883 epochs.

Table I summarizes the findings when comparing these three methods. As anticipated, MAPT outperformed standalone NavQ in every respect, demonstrating a substantially lower maximum error, significantly fewer epochs with a horizontal error exceeding 50 m, and the ability to produce a position solution for all epochs.

TABLE I
SUMMARY OF PERFORMANCE RESULTS FOR THE DATA COLLECTED IN COLUMBUS, OHIO. EWM50: EPOCHS WITH MORE THAN 50 M OF ERROR, EWNS: EPOCHS WITH NO SOLUTION. TOTAL NUMBER OF EPOCHS: 5475.

Method	**95th Percentile of Error [m]**	**Max Error [m]**	**EWM50 [%]**	**EWNS [%]**
All Meas.	47.84	206.14	3.98	0.55
NavQ	7.83	136.92	0.29	1.22
MAPT	4.44	64.76	0.07	0

The results from MAPT analyzing up to seven satellites shown in Table I are similar to the ones obtained in [4] when the proposed algorithm analyzed ten satellites. This shows a clear improvement in the algorithm's ability to select satellites more effectively using the NavQ multipath model with the NLOS masking, rather than simply analyzing the lowest elevation satellites. Additionally, the improvements to the performance efficiency of MAPT resulted in an average computation time of 0.23 s per epoch compared to 6.25 s achieved in [4]. These results clearly show an improvement when using our new methodology over prior approaches.

Table II details the execution time of MAPT's stages. The system used in this study had two Intel Xeon Silver 4216 CPUs, each with 16 cores/32 threads (hyperthreading enabled), 128 GB of DDR4 ECC RAM and an A6000 NVIDIA GPU. On average, all stages operate at real-time speeds. The total epoch processing time is the average across all epochs, including those which only executed the NavQ stage due to a high confidence computed by the model.

TABLE II
SUMMARY OF THE EXECUTION TIME FOR EACH STAGE IN MAPT ALGORITHM FOR THE 5475 EPOCHS WORTH OF DATA COLLECTED IN COLUMBUS, OHIO USING A NAVQ CONFIDENCE THRESHOLD OF 0.33.

Stage	**Mean execution time [ms]**	**Maximum execution time [ms]**
NavQ Execution	86.63	184.94
NavQ Multipath Model Prediction	163.6	210.04
Satellite Combinations and Position Computation	163.47	540.22
Clustering Algorithm	149.44	774.29
Total Epoch Processing	**232.71**	**1,360.8**

Additionally, there are certain scenarios where the complexity of the satellite combinations and position computation stage increases significantly due to the need to adaptively modify the selected satellite configuration by incorporating NLOS satellites when insufficient valid position solutions are generated for clustering. In these situations, characterized by low LOS satellite counts and high uncertainty, determining optimal satellite selection can be challenging. In the clustering algorithm stage, finding the correct value for ϵ can take longer depending on the distribution of the position solutions. The maximum value encountered for the clustering algorithm stage was an exception and all other epochs processed this stage well under 300 ms.

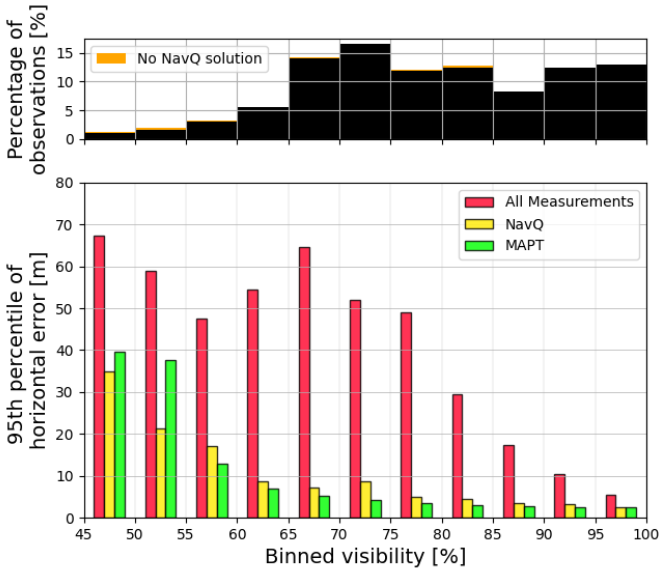

Fig. 13. 95^{th} percentile of horizontal error plotted as a function of the percentage of visibility of the sky for the three methods.

Figure 13 shows each method's 95^{th} percentile of horizontal

error as a function of the visible sky percentage where bin size is 5%. The percentage of the total observations that fell within that percentage of visibility range and the percentage of NavQ trials that were unable to produce a position solution can be observed on the top graph. The general trend observed here is that, at low percentage of sky visibility, the error when using all the measurements is large but starts improving when the visibility is 80% or higher. Both NavQ and MAPT offer a significant improvement at lower visibility ranges, and MAPT consistently produced better position solutions than NavQ above 55% visibility.

For the epochs where the visibility ranged between 45% and 55%, the results are misleading. The 95^{th} percentile of error for NavQ seems to be lower than MAPT. This discrepancy is attributed to 19% of the epochs in this visibility range where NavQ failed to produce a position solution, resulting in these epochs being excluded from the 95^{th} percentile calculation.

V. Discussion

To illustrate the effectiveness of our proposed algorithm, two specific scenarios describing how the algorithm was employed are presented. The following examples highlight both successful and challenging applications of the algorithm.

A. Successful Application

The MAPT algorithm demonstrated excellent performance in scenarios where it was evident that a dense cluster of position solutions was found and other position solutions were spread apart. For the example shown in Figure 14, a limited sky visibility can be observed in the Skyplot (a), and the satellites with the largest potential to contain multipath (b).

In this example, the sensor observed a total of 20 signals, of which ten are LOS satellites. The multipath model suggested satellites G19 and E24 were likely to have a multipath error, so the third stage of MAPT processing included these satellites in the list used to produce the combinations: G19, E24, C26, R12, G17, E33, and G14; where C44, E31 and R13 were retained as the keystone satellites due to their high elevation angles. 53 position solutions were valid, which meant that MAPT targeted a cluster with 12 positions. As shown in Figure 14(c), the selected position solution resulted in a significantly lower error. This position solution was obtained by removing satellites G19 and G17 from the list of satellites, indicating that G19 was indeed erroneous. When using all of the measurements available to the sensor, the horizontal error was 22.1 m; with NavQ removing NLOS satellites, the error increased to 117.1 m but with MAPT, the resulting horizontal error was 7.7 m.

B. Limitations of MAPT

While the proposed algorithm demonstrated impressive results in most scenarios, it was not without its limitations. Epochs within areas of dense urban canyon are challenging for the clustering algorithm due to reduced satellite visibility and poor geometry. In certain extreme circumstances, our algorithm is unable to produce enough valid combinations

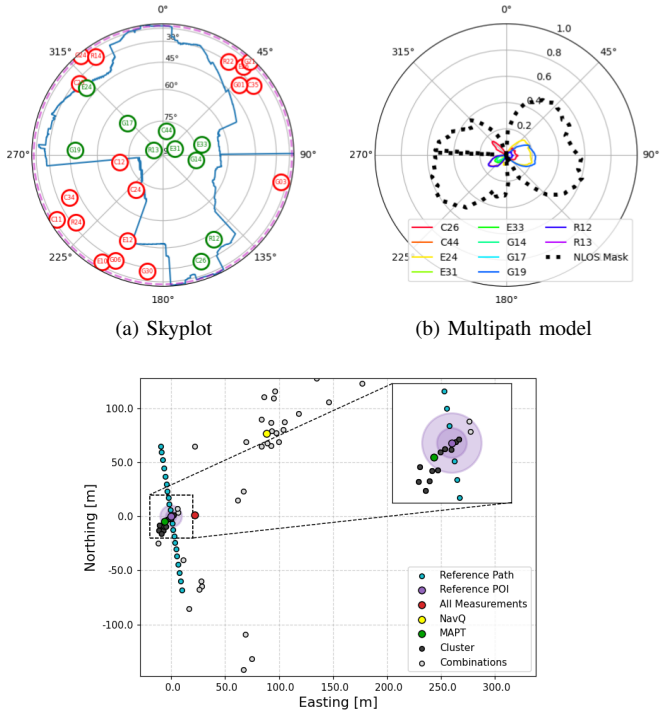

(a) Skyplot (b) Multipath model

(c) Position solutions

Fig. 14. Example outputs from MAPT for the measurements obtained by the receiver at 2023-05-31 14:51:46 UTC.

to obtain enough information to use the clustering algorithm, therefor resulting in the use of NLOS satellites to provide some solution.

Figure 15 shows an example where no method is able to find a reasonable solution. For this epoch, the receiver observed 19 signals from satellites, but only seven signals came from LOS satellites due to a 50.3% sky visibility. When MAPT processed this scenario, it was unable to obtain enough position solutions until it included four NLOS satellites, which also resulted in a long execution time for the third stage of the algorithm. The horizontal error obtained using all of the available measurements was 39.5 m. When the NLOS satellites are removed (NavQ's solution), the error was 60.6 m, and MAPT's selected solution resulted in 64.8 m of error.

The limitations of the algorithm in this scenario highlight the importance of careful consideration of environmental factors and the need for further research into improving its robustness.

C. Significance of Work

This research tackles a critical challenge in positioning technology, that is, inaccurate position estimation in urban environments using GNSS. While methods like RAIM and augmentation systems have proven effective in open sky conditions with minimal erroneous satellite measurements, new solutions are necessary to address the unique difficulties posed by low altitude flight or surface operations in urban settings. This study presents a novel approach that explores a wide

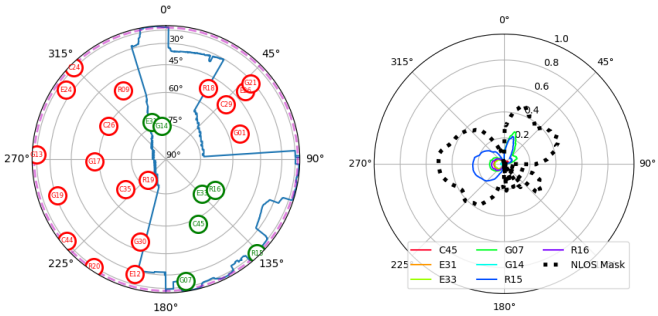

(a) Skyplot (b) Multipath model

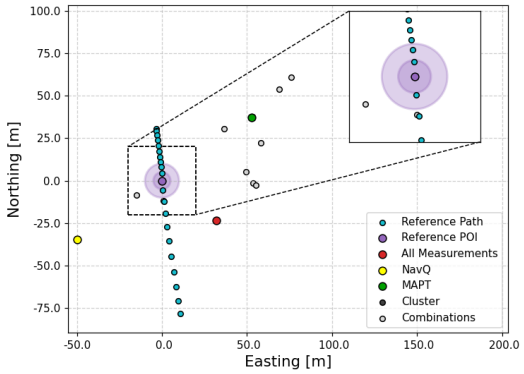

(c) Position solutions

Fig. 15. Example outputs from MAPT for the measurements obtained by the receiver at 2023-05-11 15:25:20 UTC.

range of potential satellite measurement combinations, employing advanced clustering algorithms and high performance computing methods to mitigate multipath interference.

D. Future Work

Future work will focus on developing a prototype of the framework capable of real-time execution on a local platform with the sensor. Additionally, performance enhancement by coupling this improved position solution with an inertial measurement unit (IMU) for even better accuracy will be explored. Furthermore, the identification of optimal satellite combinations that yield the lowest horizontal error and 3D positioning error using weighted clustering techniques will be examined.

VI. Conclusions

MAPT provides a robust approach for improving position estimation accuracy when significant multipath is present on multiple satellite measurements. By utilizing machine learning with the NavQ GNSS performance monitor, we introduce a novel framework that enhances position estimates, delivering a more reliable and precise navigation and localization system. The solution incorporates a receiver-agnostic, unsupervised machine learning algorithm that operates independently of specific receiver characteristics. Additionally, the multipath modeling approach avoids the complexities of modeling signal propagation, destructive and constructive interference, material properties, and other physical attributes of the environment. The effectiveness of this novel framework in two distinct scenarios was demonstrated, showcasing its ability to reduce errors in real urban environments as well as simulated multipath environments. Correct satellite selection has been demonstrated to be a key factor for the satellite combinations to allow MAPT to eliminate the faulty satellite measurements.

Acknowledgments

This work has been partially supported by the National Science Foundation under Awards ECCS-1845833 and CCF-2326559. The authors would like to thank Dr. Steve Langel (MITRE) for his valuable discussions and support during this research. The authors would also like to thank Dr. Sarah Lehman (NASA) and Dr. Steve Young (NASA) for their valuable technical feedback.

References

[1] S. Young, E. Ancel, A. Moore, E. Dill, C. Quach, J. Foster, K. Darafsheh, K. Smalling, S. Vazquez, E. Evans et al., "Architecture and information requirements to assess and predict flight safety risks during highly autonomous urban flight operations," NASA Langley Research Center, Tech. Rep., 2020.

[2] S. Hewitson and J. Wang, "GNSS receiver autonomous integrity monitoring (RAIM) performance analysis," Gps Solutions, vol. 10, 2006, pp. 155–170.

[3] A. Angrisano, S. Gaglione, and C. Gioia, "Raim algorithms for aided gnss in urban scenario," in 2012 Ubiquitous Positioning, Indoor Navigation, and Location Based Service (UPINLBS). IEEE, 2012, pp. 1–9.

[4] J. Gutierrez, R. Gilabert, E. Dill, G. Hernandez, D. Kaeli, and P. Closas, "Multipath mitigation via clustering for position estimation refinement in urban environments," in Proceedings of the ION 2024 Pacific PNT Meeting, 2024, pp. 556–568.

[5] E. Dill, J. Gutierrez, S. Young, A. Moore, A. Scholz, E. Bates, K. Schmitt, and J. Doughty, "A predictive gnss performance monitor for autonomous air vehicles in urban environments," in Proceedings of the 34th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2021), 2021, pp. 125–137.

[6] J. Gutierrez, S. Young, A. Moore, R. Gilabert, E. Dill, E. Bates, M. Peretic, K. Schmitt, and A. Scholz, "A high-performance computing predictive gnss performance monitor for autonomous air vehicles in urban environments," NASA Langley Research Center, Tech. Rep., 2024.

[7] T. Takasu, N. Kubo, and A. Yasuda, "Development, evaluation and application of rtklib: A program library for rtk-gps," in GPS/GNSS symposium, 2007, pp. 213–218.

[8] R. Gilabert, J. Gutierrez, and E. Dill, "Improvements to GNSS Positioning in Challenging Environments by 3DMA Lidar Informed Selective Satellites Usage," in Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023), 2023, pp. 2606–2615.

[9] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," Journal of Machine Learning Research, vol. 12, 2011, pp. 2825–2830.

[10] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. VanderPlas, A. Joly, B. Holt, and G. Varoquaux, "API design for machine learning software: experiences from the scikit-learn project," in ECML PKDD Workshop: Languages for Data Mining and Machine Learning, 2013, pp. 108–122.

[11] B. Hofmann-Wellenhof, H. Lichtenegger, and E. Wasle, GNSS–global navigation satellite systems: GPS, GLONASS, Galileo, and more. Springer Science & Business Media, 2007.

[12] D. Xu and Y. Tian, "A comprehensive survey of clustering algorithms," Annals of Data Science, vol. 2, 2015, pp. 165–193.

[13] N. NGS, "Guidelines for new and existing continuously operating reference stations (cors)," Document Number, 2013.