Event-Driven Geolocation Validation for Rural Autonomous Systems: TrustGeoScore and Geolocation Name System

Juan A. Padilla
Computer Science and Engineering Department
University of Puerto Rico
Mayagüez, Puerto Rico, USA
e-mail: juan.padilla11@upr.edu

***Abstract*—Rural autonomous systems suffer from unreliable geolocation due to incomplete mapping, ambiguous civic addressing, Global Navigation Satellite System (GNSS) drift, and sparse digital infrastructure. These limitations undermine safety and operational reliability for drones, agricultural automation, emergency response, and infrastructure inspection. This paper introduces TrustGeoScore, an event-driven geolocation reliability model, and Geolocation Name System (GNS), a contextual namespace framework for operationally meaningful sub-locations. TrustGeoScore aggregates multi-source evidence—including drone telemetry, Internet of Things (IoT) signals, GNSS stability observations, identity-validated user confirmations, and application-level operational validations—into a dynamic trust value with temporal decay. We present the system architecture, justify the mathematical formulation and calibration methodology, demonstrate how trust integrates into classical routing via a Dijkstra extension, present empirical structural validation using large-scale rural geolocation data, and explicitly discuss system limitations. The proposed framework enables safer and more reliable autonomous operations in rural environments.**

***Keywords*-Geolocation Reliability; Event-Driven Validation; Autonomous Systems; Trust Modeling.**

I. Introduction

Autonomous systems are now being deployed in rural regions for logistics, agriculture, emergency response, and infrastructure monitoring. However, these environments expose fundamental weaknesses in existing geolocation systems. Traditional maps and civic addresses collapse large, multifunctional properties into single coordinates, ignoring entrances, landing zones, and operational constraints. Global Navigation Satellite System (GNSS) drift further degrades precision due to terrain, vegetation, and sparse correction infrastructure [1]–[4].

In many rural contexts, autonomous systems are expected to operate with the same reliability as in urban environments, despite receiving significantly poorer geospatial inputs. This mismatch creates safety risks, increases operational costs, and limits adoption. Addressing this challenge requires rethinking geolocation not as a static coordinate, but as a dynamic, evidence-backed construct whose reliability can be measured and reasoned about.

Prior work has explored individual components relevant to geolocation reliability, including alternative addressing schemes, GNSS integrity monitoring, and navigation under uncertainty [1][2][4]–[6]. However, these efforts address isolated dimensions of the problem and do not integrate contextual namespaces, event-driven validation, temporal trust decay, and application-level operational confirmation within a single framework. The present work addresses this gap by unifying these dimensions into a cohesive, deployable system.

This paper addresses these challenges by introducing TrustGeoScore and Geolocation Name System (GNS). Together, they provide a contextual, event-driven geolocation framework designed explicitly for rural autonomous operations.

The remainder of this paper is structured as follows. Section II reviews background and motivates the need for dynamic geolocation reliability modeling. Section III introduces GNS. Section IV presents the TrustGeoScore model, including aggregation and decay mechanisms. Section V describes the system architecture. Section VI illustrates practical applications. Section VII presents the empirical structural validation. Section VIII analyzes design trade-offs and limitations, and Section IX concludes the paper.

II. Background and Motivation

Geolocation systems were designed primarily for human navigation. Humans can interpret imprecise directions, identify landmarks, and adapt to uncertainty; autonomous systems cannot. Rural environments amplify this mismatch due to incomplete or outdated maps, ambiguous civic addresses, GNSS drift and multipath interference, and sparse digital infrastructure [2][4][5][7][8]. Recent studies further highlight the vulnerability of GNSS signals to spoofing, interference, and signal integrity degradation in complex environments [9]. These limitations are consistent with longstanding observations in geographic information science that human-generated spatial data and addressing systems are inherently uneven and context-dependent, particularly outside dense urban areas [10].

Puerto Rico provides a representative example of broader rural geolocation challenges. U.S. Census assessments document widespread limitations in civic address coverage, where many residences lack standardized, consistently usable street addresses and instead rely on informal or descriptive location references [11]. Similar challenges in geocoding and address standardization have been documented in broader urban and developing contexts, where heterogeneous and incomplete address data can undermine spatial reliability [12]. Such variability in address-to-coordinate translation further complicates autonomous system operation when geolocation inputs are treated as authoritative. Comparable conditions exist across

 ISBN: 978-1-68558-366-8

many rural regions in the continental United States, where properties span large areas, contain multiple access points, and are poorly captured by digital maps.

As autonomy expands into rural logistics, precision agriculture, emergency response, and infrastructure inspection, geolocation reliability becomes a limiting factor. A new paradigm is required—one that models reliability dynamically rather than assuming static correctness.

While this work focuses on rural environments, similar shortcomings appear in dense urban areas and developing regions where informal addressing, rapid construction, or incomplete municipal digitization create inconsistencies between physical reality and digital maps. In many developing countries, buildings may lack standardized addresses, entrances may be unmarked, and GNSS multipath effects degrade reliability in narrow streets. In such contexts, TrustGeoScore can complement formal addressing schemes by incorporating operational validation signals, enabling autonomous systems to reason about reliability even when official records are incomplete or outdated.

III. Geolocation Name System

GNS introduces contextual geolocation namespaces that represent operationally distinct sub-locations within a property. Instead of mapping an entire property to a single coordinate, GNS allows multiple semantic identifiers such as:

- `farm.main-entrance`
- `farm.drone-pad.south`
- `farm.irrigation-zone.3`
- `home.dropoff.backyard`

Each namespace maps to coordinates, metadata, access permissions, and a TrustGeoScore value. This enables autonomous systems to reason about what a location represents and how reliable it is, rather than relying on a single ambiguous point.

Namespaces capture operational semantics that raw coordinates cannot express, such as preferred approach direction, obstacle constraints, surface type, and mission suitability. This is particularly important in rural environments where multiple operational sub-locations may exist within a single property boundary.

IV. TrustGeoScore Model

TrustGeoScore assigns a dynamic reliability score to each GNS namespace based on real-world evidence.

A. Evidence Sources

TrustGeoScore aggregates evidence, including:

- drone landings and flight telemetry,
- Internet of Things (IoT) sensor pings from fixed infrastructure,
- GNSS stability and drift observations,
- identity-validated human confirmations,
- application-level operational validations, such as delivery confirmations, inspection completions, or service acknowledgments.

Application-level validations are particularly important because they represent end-to-end confirmation that a geolocation successfully supported a real-world operation. A completed delivery or confirmed service task provides strong evidence that the location was not only reachable, but operationally correct.

B. Aggregation and Decay

Trust is computed using weighted aggregation:

$$TG(v,t) = \frac{\sum_{k=1}^{m} w_k \, TG_k(v,t)}{\sum_{k=1}^{m} w_k} \tag{1}$$

and temporal decay:

$$TG(v,t) = TG(v,t_0)\, e^{-\lambda(t-t_0)} + \mathrm{NewEvidence}(t) \tag{2}$$

where v represents the GNS namespace (location identifier); t is the current time (seconds or any consistent time unit); $TG(v,t)$ is the TrustGeoScore of namespace v at time t, a value in $[0,1]$; m is the total number of evidence source types; w_k is the non-negative weight assigned to source k (calibrated per source reliability); $TG_k(v,t)$ is the trust subscore contributed by source k, in $[0,1]$; t_0 is the time of the most recent trust update (seconds); $\lambda > 0$ is the decay rate constant (s^{-1}, controlling the half-life $t_{1/2} = \ln 2/\lambda$); and $\mathrm{NewEvidence}(t)$ is the incremental trust contribution from validation events arriving at time t, in $[0,1]$.

Weighted linear fusion is chosen for interpretability and robustness under missing data. Weights (w) represent expected source reliability and are calibrated via source characterization and optional data-driven optimization. Contemporary trust management frameworks in IoT and cyber-physical systems similarly emphasize lightweight, interpretable scoring mechanisms suitable for distributed environments [13]. Exponential decay is selected due to its standard use in modeling information staleness and its intuitive half-life interpretation [14]. All parameters are explicitly defined to ensure transparency and reproducibility.

Exponential temporal decay was adopted to model information staleness in a manner consistent with both human intuition and established practices in trust and information reliability modeling. By gradually reducing the influence of outdated evidence rather than imposing hard expiration thresholds, the model avoids abrupt trust discontinuities while still reflecting growing uncertainty over time.

C. Trust Evolution and Interpretation

TrustGeoScore is designed to represent geolocation reliability as a continuous, evidence-driven quantity rather than a static or binary property. This section clarifies the intended semantic meaning of the trust score and the role of its primary components—aggregation, decay, and evidence weighting—without prescribing runtime behavior or decision logic.

At a conceptual level, TrustGeoScore captures the degree of confidence that a given GNS namespace will successfully support real-world operations under its intended context. Rather than expressing absolute positional accuracy, the score reflects operational reliability, integrating heterogeneous validation signals into a single interpretable measure. Reputation-based and integrity-aware trust mechanisms have been widely studied in sensor networks and distributed systems [15][16], supporting the use of weighted evidence aggregation in safety-critical environments.

By modeling trust as an evolving belief rather than a definitive guarantee, TrustGeoScore provides a semantic abstraction that can be consumed uniformly by autonomous systems, human-in-the-loop workflows, and downstream applications.

D. Role of Application-Level Validation

A distinguishing feature of TrustGeoScore is its explicit incorporation of application-level operational validation as a high-value source of trust evidence. While sensor data and GNSS observations provide important signals about spatial consistency, they do not by themselves confirm that a geolocation successfully supported a complete real-world operation. Application workflows, by contrast, provide end-to-end confirmation that a location was operationally correct.

Examples of application-level validation include successful package delivery confirmations, completed inspection tasks, verified agricultural operations, or acknowledged emergency supply drops. When an application confirms task completion, the system infers that the geolocation used for that operation was reachable, safe, and contextually appropriate. Such confirmations implicitly validate multiple dimensions at once: accessibility, spatial accuracy, environmental suitability, and alignment with operational intent.

In the TrustGeoScore framework, application-level validations are treated as particularly strong evidence when they are identity-validated and associated with a well-defined task. A confirmed delivery, for instance, reinforces trust more effectively than a single GNSS reading because it demonstrates that the geolocation supported a successful outcome under real operating conditions. Conversely, repeated task failures or aborted operations contribute negative evidence, reducing trust and signaling potential issues such as drift, obstruction, or misclassification of the location.

This approach enables a closed-loop validation process in which geolocations are continuously refined through normal system usage rather than through dedicated calibration efforts. As autonomous systems operate, application outcomes feed directly back into trust computation, allowing the system to learn which locations consistently work and which do not. This closed-loop behavior is particularly well suited to rural environments, where explicit ground truth is scarce but operational feedback is naturally generated through routine activities.

By elevating application-level outcomes to first-class trust signals, TrustGeoScore bridges the gap between abstract geospatial correctness and practical operational reliability.

E. Interpretation of TrustGeoScore Dynamics

While the preceding subsection defines the semantic meaning of TrustGeoScore, this section describes how the trust value behaves over time as new evidence arrives, locations are repeatedly used, or validation activity diminishes. These dynamics are essential for enabling safe and adaptive decision-making in autonomous and semi-autonomous systems.

In practice, TrustGeoScore evolves through three characteristic phases:

- *Initialization Phase* — Newly created GNS namespaces begin with a conservative baseline trust value. During this phase, individual validation events may exert a relatively strong influence on the trust score, reflecting uncertainty due to limited accumulated evidence. This conservative initialization discourages premature reliance on unproven geolocations.
- *Stabilization Phase* — As consistent validation events accumulate—such as repeated drone landings, IoT sensor confirmations, or application-level task completions—the trust score converges toward a stable value. Event diversity plays a critical role during this phase, as corroboration from independent sources increases confidence more robustly than repeated confirmation from a single signal type.
- *Decay and Revalidation Phase* — When a namespace is not exercised for extended periods, temporal decay gradually reduces trust, reflecting uncertainty about whether environmental conditions or access constraints may have changed. Subsequent validation restores trust incrementally rather than instantaneously, preventing abrupt trust spikes based on isolated confirmations.

Through these dynamics, TrustGeoScore supports cautious reliance on geolocations whose reliability is continuously reinforced through use, while naturally reducing confidence in locations that become stale or insufficiently validated. This behavior ensures that trust reflects current operational reliability rather than historical correctness alone.

These trust dynamics enable downstream systems to reason about geolocation reliability in a principled and adaptive manner, but their practical impact depends on how they are realized within a deployable system. The following section describes the system architecture that operationalizes TrustGeoScore within the GNS framework, detailing the components and data flows that support evidence ingestion, trust computation, and integration with autonomous and human-in-the-loop applications.

V. System Architecture

The proposed system architecture operationalizes TrustGeoScore within GNS as a modular, event-driven framework designed to support trust-aware geolocation management under real-world uncertainty. The architecture ingests heterogeneous validation evidence, enforces identity and authorization constraints, computes dynamic trust scores, and exposes trust-aware services to autonomous and human-in-the-loop applications.

To support deployment in environments with intermittent connectivity and heterogeneous infrastructure, the architecture emphasizes loose coupling, incremental updates, and graceful degradation. Trust quality improves progressively as additional validation events are observed, allowing the system to function effectively even during early or sparse deployment stages.

A. Core Components

The system consists of eight primary components, each responsible for a distinct stage in the trust lifecycle:

- *Event Collector* — Ingests validation events originating from drones, IoT devices, application workflows, and human users. Events are captured opportunistically during normal operations, minimizing the need for dedicated validation procedures.
- *Event Normalizer* — Transforms heterogeneous event formats into a standardized representation that preserves source identity, temporal context, and validation outcome, enabling uniform downstream processing.
- *Authorization and Identity Validation Engine* — Verifies the authenticity of actors generating events and enforces access permissions, enabling differentiated trust weighting while supporting accountability and controlled evidence contribution.
- *Validation Engine* — Evaluates spatial, temporal, and contextual consistency of incoming events relative to the referenced GNS namespace. This step filters inconsistent or implausible evidence before it influences trust computation.
- *Trust Engine* — Computes TrustGeoScore values using the weighted aggregation and temporal decay mechanisms described in Section IV. Trust updates are performed incrementally as new evidence arrives, enabling continuous adaptation.
- *GNS Registry* — Maintains persistent geolocation namespaces along with associated metadata and historical trust trajectories. This registry decouples geolocation identity from transient sensing conditions.
- *API Layer* — Exposes trust-aware geolocation services to external applications, allowing systems to query locations together with their associated reliability metrics without direct access to raw validation events.
- *Analytics Dashboard* — Provides visualization of trust evolution, evidence contributions, and uncertainty trends to support monitoring, auditing, and human oversight.

Figure 1 illustrates the event-driven data flow and cross-stage observability within the proposed architecture.

B. Design Considerations

The architecture is explicitly designed to accommodate uncertain, evolving environments where ground truth is incomplete or unavailable. Incremental trust computation and temporal decay allow the system to adapt naturally to changing conditions without requiring manual reconfiguration. Privacy and accountability are supported through identity validation and controlled evidence exposure, while modularity enables selective deployment of components based on operational needs.

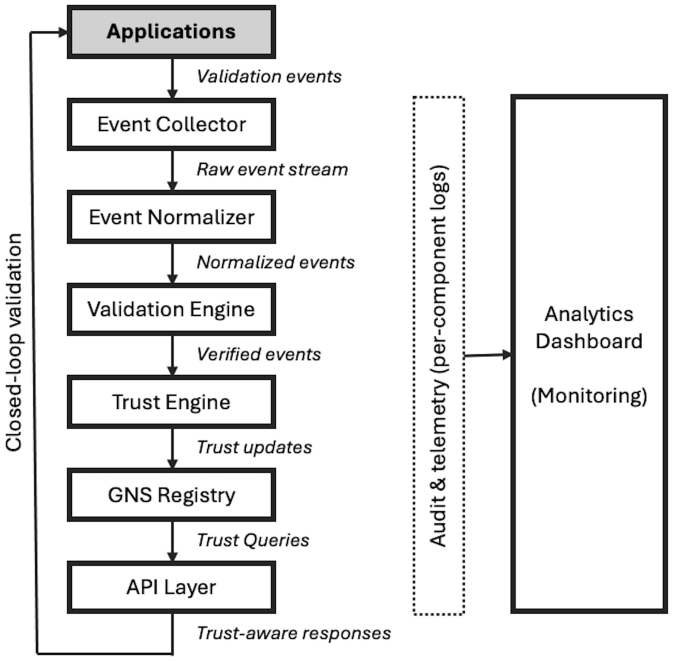

Figure 1. Event-driven TrustGeoScore architecture. Validation events are collected, normalized, verified, aggregated into trust scores, and persisted in the GNS Registry, with per-component audit and telemetry logs supporting monitoring.

C. Operational Lifecycle and Deployment

Deploying TrustGeoScore and GNS in real rural environments requires consideration of the full operational lifecycle of geolocations, from initial creation through long-term maintenance and retirement. Unlike static mapping systems, which implicitly assume that locations remain valid indefinitely, TrustGeoScore treats geolocations as evolving entities whose reliability must be continuously reassessed.

1) Geolocation Initialization and Bootstrapping: New GNS namespaces are typically created through a combination of user input, application workflows, and automated discovery processes. During initialization, namespaces may be populated with approximate coordinates, descriptive metadata, and initial access permissions. At this stage, TrustGeoScore assigns a conservative baseline trust value, reflecting the absence of empirical validation. This design choice prioritizes safety by discouraging autonomous systems from relying heavily on newly defined locations until sufficient evidence has accumulated.

Bootstrapping may be accelerated through contextual priors, such as proximity to existing high-trust namespaces or similarity to previously validated operational patterns. However, these priors are deliberately limited in influence to avoid masking genuine uncertainty.

2) Long-Term Maintenance and Drift Management: Over time, environmental conditions change. Roads erode, vegetation grows, access points shift, and infrastructure is modified.

TrustGeoScore accounts for these dynamics through temporal decay, which gradually reduces confidence in locations that are not revalidated. This mechanism prevents outdated geolocations from remaining falsely trusted and encourages periodic revalidation.

In practice, this means that frequently used and operationally stable locations maintain high trust, while rarely used or abandoned locations naturally decline in reliability.

VI. Applications

This section illustrates how TrustGeoScore and GNS are applied in practice to support reliable autonomous operation in rural and infrastructure-sparse environments. Rather than treating geolocation as a static input, the proposed framework enables systems to adapt behavior based on continuously updated trust signals derived from real-world operational evidence.

A. Trust-Aware Drone Operations in Rural Environments

Rural properties often contain multiple entrances, landing zones, and staging areas, yet are typically represented by a single civic address or coordinate. Autonomous systems navigating to such locations may encounter obstacles or unsafe conditions, particularly when GNSS drift shifts approach points by several meters.

Using GNS, a property can define multiple contextual namespaces representing operationally distinct sub-locations, such as landing pads, access roads, or delivery zones. Each namespace is associated with coordinates, metadata, and a TrustGeoScore value that reflects accumulated validation evidence. Prior to execution, an autonomous drone queries the system to identify candidate namespaces and their associated trust scores.

High-trust namespaces indicate locations that have consistently supported successful operations, allowing missions to proceed autonomously with minimal additional verification. When trust values are moderate or declining, the system adapts behavior rather than failing outright. For example, the drone may adopt a more conservative approach strategy, request additional sensing confirmation, or escalate to human oversight before committing to landing.

As missions complete, application-level outcomes feed directly back into the trust computation process. Successful landings, confirmed deliveries, or completed inspections incrementally reinforce trust, while aborted missions or access failures contribute negative evidence. Over time, this closed-loop feedback allows the system to learn which rural sub-locations reliably support operations and which degrade due to environmental change, obstruction, or misclassification.

B. Trust-Aware Routing and Navigation

TrustGeoScore integrates directly into classical routing algorithms by treating geolocation reliability as a first-class routing parameter. In the case of Dijkstra's algorithm [17], the relaxation step is modified to incorporate a trust-based penalty:

$$d[v] = \min\big(d[v],\ d[u] + w(u,v) + g\big(TG(v)\big)\big) \qquad (3)$$

where $d[v]$ is the tentative shortest-path cost to node v (meters or cost units); $d[u]$ is the settled cost to the current node u; $w(u,v)$ is the nominal edge weight between nodes u and v (meters or cost units); $TG(v)$ is the TrustGeoScore of node v, in $[0,1]$; and $g(\cdot)$ is a monotonically decreasing penalty function that maps trust to an additive cost (cost units), penalizing low-trust nodes with a larger routing cost.

Trust-aware routing complements existing link-state routing approaches commonly used in distributed and ad hoc networks [18], while preserving the structure of classical shortest-path algorithms.

Low-trust nodes incur higher penalties, steering routes away from unreliable or stale locations while preserving the underlying algorithmic structure. This approach improves navigation robustness without requiring custom routing algorithms or abandoning established graph-based methods.

By integrating trust into routing decisions, autonomous systems can balance distance, cost, and reliability simultaneously. Routes that are slightly longer but more trustworthy may be preferred over shorter paths that terminate at uncertain or poorly validated locations, particularly in safety-critical or time-sensitive operations.

C. Domain-Specific Impact

Projected benefits of trust-aware geolocation are consistent with reported challenges in rural autonomy and infrastructure operations [4][6][19]. Conservative estimates derived from prior failure rates and pilot-scale observations suggest the following improvements:

- *Logistics:* 40–60% reduction in landing misidentification and 15–25% fewer mission aborts.
- *Agriculture:* 25–45% improvement in operational consistency and 10–20% improved sampling repeatability.
- *Emergency Response:* 50–70% reduction in navigation failures and 30–50% faster arrival times.
- *Telecommunications:* 15–25% reduction in inspection time and fewer failed approach vectors.

These figures represent projected improvements rather than results from large-scale deployments. They are not measured outcomes from the empirical analysis presented in Section VII. They serve to illustrate the practical impact of incorporating dynamic geolocation trust into autonomous workflows under rural operating conditions.

D. Closed-Loop Validation via App-Level Events

In deployed applications, TrustGeoScore enables a closed-loop validation process in which operational outcomes feed directly back into geolocation trust. Unlike passive sensing, application workflows provide end-to-end confirmation that a geolocation supported a successful real-world task.

Examples of such application-level events include:

- successful package delivery confirmations.
- completed agricultural missions.
- inspection task completions.
- emergency supply drop acknowledgments.

TABLE I. SPATIAL METRICS OF INSTANTIATED GNS NAMESPACES

Metric	Value
Total unique geolocation IDs	32,019
Geographic coverage area	199.7 km^2
Namespace density	160.4 per km^2
Median nearest-neighbor distance	4.5 m
Median distance to any mapped road	10.4 m
Median distance to major road	179.1 m
75th percentile distance to major road	323.7 m
95th percentile distance to major road	580.9 m

When an application confirms task completion, the system infers that the geolocation used for that operation was not only reachable, but operationally correct. These events are particularly valuable because they implicitly validate multiple dimensions at once: accessibility, safety, spatial accuracy, and contextual suitability.

In the TrustGeoScore framework, application-level confirmations are treated as high-confidence evidence, especially when identity-validated (e.g., authenticated users, delivery systems, or enterprise applications). This creates a closed feedback loop:

1) GNS provides a contextual geolocation.
2) TrustGeoScore selects a high-trust candidate.
3) The autonomous system executes the task.
4) The application confirms success.
5) TrustGeoScore is reinforced.

This loop enables continuous improvement of geolocation reliability without manual intervention and distinguishes TrustGeoScore from purely sensor-based validation approaches.

VII. EMPIRICAL STRUCTURAL VALIDATION

To empirically ground the proposed framework, we conducted a large-scale spatial analysis using 32,019 unique geolocation records instantiated under a United States Department of Agriculture supported delivery initiative in central Puerto Rico. These records were used to bootstrap GNS identifiers, assigning self-validating namespaces to operationally meaningful locations. Although repeated deployment phases were not completed due to program discontinuation, the dataset provides a substantial basis for evaluating spatial structure, infrastructure proximity, and namespace scalability.

A. Dataset Scale and Coverage

As summarized in Table I, the instantiated namespaces span approximately 199.7 km^2, with a density of 160.4 locations per km^2. The region includes dense urban neighborhoods, suburban developments, and peri-urban and rural zones.

The median nearest-neighbor distance between instantiated namespaces is 4.5 meters, indicating localized clustering consistent with residential environments and multi-unit structures. This clustering coexists with broader regional dispersion.

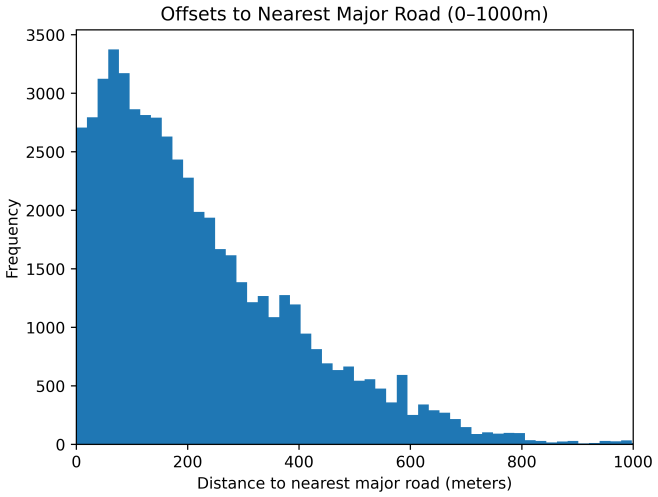

Figure 2. Distribution of distances from instantiated GNS namespaces to the nearest major road infrastructure (motorway, trunk, primary, secondary, tertiary). The right-skewed distribution highlights structural variability in infrastructure proximity across the study region.

B. Proximity to Road Infrastructure

To evaluate the relationship between instantiated namespaces and transportation infrastructure, distances were computed to both (1) any mapped roadway and (2) major road infrastructure (motorway, trunk, primary, secondary, and tertiary classifications from OpenStreetMap).

As shown in Table I, the median distance to any mapped road is 10.4 meters, indicating that most locations are adjacent to some form of roadway. However, the median distance to the nearest major road is 179.1 meters. Furthermore, 75% of locations lie within 323.7 meters of a major road, while 5% exceed 580.9 meters.

The full distribution of offsets to major road infrastructure is illustrated in Figure 2. The right-skewed distribution demonstrates that while many locations are within several hundred meters of major corridors, a substantial portion are located deeper within residential or secondary road networks. Prior analyses of road network hierarchy demonstrate that accessibility and mobility patterns are strongly shaped by hierarchical street classifications, particularly in mixed urban–rural environments [20].

These structural offsets motivate the trust weighting mechanisms defined in Section IV, particularly in scenarios where proximity to primary infrastructure cannot be assumed.

C. Implications for Contextual Geolocation Modeling

These findings reveal a structural distinction between proximity to any roadway and proximity to primary infrastructure. While most geolocations are near some mapped road, operational access frequently depends on minor or residential networks rather than major transportation corridors.

This stratified access structure reinforces the need for contextual namespace abstraction (GNS) and event-driven reliability estimation (TrustGeoScore), rather than assuming static

correctness based solely on civic addressing or proximity to primary roads.

D. Scope and Future Evaluation

The present analysis validates structural necessity and scalability of the proposed framework. It does not measure operational performance improvement, as repeated validation events were not available. Future deployments will evaluate routing reliability, mission success rates, and trust convergence dynamics under repeated autonomous operations.

VIII. Design Trade-offs, Alternatives, and Limitations

The design of TrustGeoScore and GNS reflects a series of deliberate trade-offs intended to balance rigor, scalability, interpretability, and practical deployability. This section discusses key alternatives and explains why the chosen approach is appropriate for rural autonomous systems.

A. Trust Modeling Versus Deterministic Validation

One alternative to trust-based modeling is deterministic validation, where a geolocation is either considered valid or invalid. While appealing in its simplicity, deterministic approaches fail to capture partial knowledge, uncertainty, and gradual degradation. Rural environments rarely provide binary certainty; access points may be usable under some conditions but not others.

TrustGeoScore instead models reliability as a continuous value, enabling autonomous systems to reason probabilistically about risk, allowing for graceful degradation rather than abrupt failure when evidence is incomplete.

B. Rule-Based Scoring Versus Learned Models

Another alternative is to use machine learning models to predict geolocation reliability directly from raw data. While potentially powerful, such approaches introduce challenges:

- dependence on large labeled datasets,
- reduced interpretability,
- difficulty generalizing across regions,
- sensitivity to data drift.

TrustGeoScore adopts a rule-based, mathematically interpretable formulation that can be calibrated using limited data and domain knowledge. This choice prioritizes transparency, reproducibility, and safety—key requirements for infrastructure and emergency applications.

C. Centralized Versus Distributed Validation

Geolocation validation could be centralized in a cloud service or distributed across edge devices. TrustGeoScore is designed to support hybrid deployment models. Core trust aggregation may occur centrally for consistency, while local validation and evidence collection can occur at the edge. This flexibility is especially important in rural environments with intermittent connectivity.

D. Manual Curation Versus Automated Evolution

Traditional geolocation systems rely heavily on manual curation and infrequent updates. While human oversight remains important, it does not scale to the complexity and dynamism of rural autonomous operations. TrustGeoScore emphasizes automated evolution driven by operational events, reducing reliance on manual intervention while preserving human override where necessary.

E. Why the Chosen Design Is Appropriate

The selected design prioritizes:

- interpretability over opaque prediction,
- gradual trust evolution over binary decisions,
- operational feedback over static assumptions,
- safety over aggressive optimization.

These trade-offs make TrustGeoScore particularly well suited for rural environments, where uncertainty is the norm rather than the exception.

F. Limitations

TrustGeoScore depends on sufficient event density; extremely sparse environments slow convergence and increase variance. New locations face cold-start challenges. Computational overhead, identity-validation latency, sensor integrity, parameter sensitivity, environmental bias, and conservative routing behavior are acknowledged limitations. These define the framework's boundaries of applicability rather than invalidating its contributions.

IX. Conclusion and Future Work

TrustGeoScore and GNS provide a contextual, event-driven geolocation validation framework tailored for rural autonomous systems. The proposed approach formally defines weighted evidence aggregation and temporal decay mechanisms, integrates trust-aware reasoning into classical routing via a Dijkstra-based extension, and operationalizes the model through a modular system architecture. Empirical structural validation using large-scale rural geolocation data demonstrates the necessity of context-aware reliability modeling in infrastructure-sparse environments. Together, these contributions establish a mathematically grounded and operationally viable foundation for safer and more reliable autonomous deployment beyond traditional static addressing schemes.

Future work will focus on full operational evaluation under repeated autonomous deployments, including routing reliability analysis, mission success rate comparisons, and empirical study of trust convergence dynamics under varying event densities. Additional investigation will explore sensitivity to parameter calibration, scalability under distributed edge deployments, and extension of the framework to dense urban and developing-region contexts with heterogeneous addressing systems.

REFERENCES

[1] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*, Rev. 2nd. Lincoln, MA, USA: Ganga-Jamuna Press, 2011, ISBN: 978-0-9709544-2-8.

[2] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd. Boston, MA, USA: Artech House, 2013, ISBN: 978-1-6080700-5-3.

[3] P. D. Groves, Z. Jiang, L. Wang, and M. K. Ziebart, "Intelligent urban positioning using multi-constellation GNSS with 3D mapping and NLOS signal detection," in *Proc. 6th ESA Workshop on Satellite Navigation Technologies (NAVITEC 2012)*, Dec. 2012, pp. 1–8. DOI: 10.1109/NAVITEC.2012.6423047.

[4] N. Zhu, J. Marais, D. Bétaille, and M. Berbineau, "GNSS position integrity in urban environments: A review of literature," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 9, pp. 2762–2778, Sep. 2018. DOI: 10.1109/TITS.2017.2766768.

[5] M. Haklay and P. Weber, "OpenStreetMap: User-generated street maps," *IEEE Pervasive Computing*, vol. 7, no. 4, pp. 12–18, 2008. DOI: 10.1109/MPRV.2008.80.

[6] Y. Chang et al., "A review of UAV autonomous navigation in GPS-denied environments," *Robotics and Autonomous Systems*, vol. 170, p. 104 533, Dec. 2023. DOI: 10.1016/j.robot.2023.104533.

[7] M. S. Braasch and A. J. Van Dierendonck, "GPS receiver architectures and measurements," *Proceedings of the IEEE*, vol. 87, no. 1, pp. 48–64, Jan. 1999. DOI: 10.1109/5.736341.

[8] Federal Communications Commission, "Wireless E911 location accuracy requirements," Federal Communications Commission, Fourth Report and Order PS Docket No. 07-114, FCC 15-9, Feb. 2015. [Online]. Available: https://www.911.gov/assets/Wireless-E911-Location-Accuracy-Requirements-1638567121.pdf.

[9] L. Meng, L. Yang, W. Yang, and L. Zhang, "A survey of GNSS spoofing and anti-spoofing technology," *Remote Sensing*, vol. 14, no. 19, p. 4826, 2022. DOI: 10.3390/rs14194826.

[10] M. F. Goodchild, "Citizens as sensors: The world of volunteered geography," *GeoJournal*, vol. 69, no. 4, pp. 211–221, 2007. DOI: 10.1007/s10708-007-9111-y.

[11] U.S. Census Bureau, *Street addresses are simple — but not in Puerto Rico*, Census.gov, Jan. 2020. [Online]. Available: https://www.census.gov/library/stories/2020/01/street-addresses-are-simple-not-in-puerto-rico.html.

[12] V. Pérez and C. Aybar, "Challenges in geocoding: An analysis of R packages and integration of external APIs for address geocoding," *ISPRS International Journal of Geo-Information*, vol. 13, no. 6, p. 170, 2024. DOI: 10.3390/ijgi13060170.

[13] M. Aaqib, A. Ali, L. Chen, and O. Nibouche, "IoT trust and reputation: A survey and taxonomy," *Journal of Cloud Computing*, vol. 12, p. 42, 2023. DOI: 10.1186/s13677-023-00416-8.

[14] Z. Yan, P. Zhang, and A. V. Vasilakos, "A survey on trust management for Internet of Things," *Journal of Network and Computer Applications*, vol. 42, pp. 120–134, Jun. 2014. DOI: 10.1016/j.jnca.2014.01.014.

[15] A. Perrig, R. Szewczyk, J. D. Tygar, V. Wen, and D. E. Culler, "SPINS: Security protocols for sensor networks," *Wireless Networks*, vol. 8, no. 5, pp. 521–534, Sep. 2002. DOI: 10.1023/A:1016598314198.

[16] S. Ganeriwal and M. B. Srivastava, "Reputation-based framework for high-integrity sensor networks," *ACM Transactions on Sensor Networks*, vol. 4, no. 3, pp. 1–37, May 2008. DOI: 10.1145/1362542.1362546.

[17] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959. DOI: 10.1007/BF01386390.

[18] T. Clausen and P. Jacquet, "Optimized link state routing protocol (OLSR)," Internet Engineering Task Force, RFC 3626, Oct. 2003. [Online]. Available: https://datatracker.ietf.org/doc/rfc3626/.

[19] Food and Agriculture Organization of the United Nations, "Digital technologies in agriculture and rural areas: Status report," FAO, Rome, Italy, Tech. Rep. CA4985EN, 2019. [Online]. Available: http://www.fao.org/3/ca4985en/ca4985en.pdf.

[20] S. Tsigdinos, A. Nikitas, and E. Bakogiannis, "Contextualizing urban road network hierarchy and its role for sustainable transport futures: A systematic literature review using bibliometric analysis and content analysis tools," *Frontiers in Engineering and Management*, vol. 12, pp. 361–393, 2025. DOI: 10.1007/s42524-024-0300-x.