A Framework for Composable Location Verification

Multi-factor location verification for credible location-contingent commitments across use cases

John Robison Hoopes

Abstract

Verifying claims about the timing and location of events is an emerging requirement across domains — from AI compute governance, where hardware-enabled mechanisms depend on knowing where advanced chips are located and operating, to autonomous systems whose permitted behavior is conditioned on geography. No single verification technique serves every setting, and each has its own vulnerabilities. We present a framework that structures location verification as a process of evidence evaluation and decision-making: *location verification systems* produce signals that are processed into signed *location stamps*; stamps compose into *location evidence* supporting a spatiotemporal *location claim*; *evidence functions* evaluate evidence against claims to produce multidimensional *credibility assessments*; and application-specific decision rules map assessments to decisions. By standardizing how evidence is structured while leaving its evaluation contextual, the framework lets independent verification systems compose — raising the cost of forgery beyond what any single system provides — while remaining neutral about which evidence types and evaluation criteria each application requires. Reliable location verification is required for credible location-contingent commitments, and foundational for the spatial governance of intelligent machines.

1. Introduction

Claims about where things happen are becoming consequential in new ways. Proposals for AI compute governance increasingly assume we can verify where advanced chips are located and operating — a prerequisite for export-control enforcement and for hardware-enabled governance mechanisms more broadly (Chip Security Act, 2025; Brass & Aarne, 2024). Autonomous systems — delivery drones, robots, vehicles — operate under rules conditioned on geography, from airspace restrictions to jurisdictional boundaries. And beyond these, any application that acts on a claim that an event occurred at a particular place and time — supply-chain audits, environmental monitoring, infrastructure networks that care about geographic distribution — inherits the same underlying problem: deciding whether to believe the claim.

No single technique serves all of these settings. Our ongoing review of location verification techniques has identified approaches ranging from physics-based measurement to cryptographic protocols to human trust networks. Each is intended to increase the difficulty of deception, and each makes different tradeoffs across security, accuracy, cost, and other dimensions. That diversity creates both challenges and opportunities:

- How do we meaningfully compare different location verification systems?
- How might signals from uncorrelated systems combine to provide stronger guarantees than any single system alone?
- How can we quantify the strength of location evidence? What dynamics impact this?
- How can developers and policymakers navigate this complexity without deep expertise in each approach?

We approach location verification through the lens of *evidence evaluation*, which provides a way to quantify confidence in evidence-based claims about where things happened. Recognizing this common pattern points toward a framework for composable location evidence — one that enables different verification systems to contribute independent evidence, while allowing verifiers to define for themselves what constitutes sufficient evidence for their needs. This makes verification both more accessible and more robust. This note presents the state of that framework.

2. Verifying Location

For such a simple idea, this is a surprisingly slippery concept.

The problem, stated plainly, is to convince a remote, skeptical verifier that a claim about the timing and location of an event is true. Concretely: this drone delivered a package at this address; this parcel crossed this border on this date; this person was not present in country Y for more than ninety days last year; this chip is operating in this jurisdiction; this sensor measurement was recorded at this site.

Key terms we're working with:

- **Location verification system** — the complete arrangement of hardware, software, and human / institutional processes that produces verifiable location evidence: one or more *signal sources*, a *stamp generator*, and the trust anchors (keys, secure hardware, institutional accountability) that make its output checkable. Sometimes shortened to "LVS".[1]

- **Location claim** — a statement that an event occurred within a specified spatial region and time interval. This often will refer to the presence of some entity, but the definition is deliberately broad.

- **Signal source** — the external infrastructure or environment an LVS observes: satellite constellations, network landmark nodes, radio environments, visual scenes, peer devices, authorities, etc.

- **Stamp generator** — the software component of an LVS that processes observed signals into a signed location stamp, applying the system's localization model. The trustworthiness of a stamp depends on where its generator runs (open software, secure enclave, tamper-resistant hardware) and on the authenticity of its inputs.

- **Signals** — raw observations from an LVS, ranging from physical measurements (RF timings, sensor data) to digital and cryptographic artifacts (identifiers, digital signatures, attestations) to assertions by external parties.[1]

- **Location stamp** — a verifiable digital artifact that corroborates a claim about the position and timing of some event, requiring collusion, technical manipulation, or fraud to forge (the term originates with Kabatnik & Zugenmaier, 2001). Produced by a stamp generator from its signals, including provenance and any relevant confidence indicators.

- **Location evidence** — a composite artifact formed by combining one or more location stamps to support a single spatiotemporal location claim.

- **Location manifest** — a location claim together with the location evidence supporting it: the pair (C, E) that a prover submits to a verifier.

- **Evidence function** — a function that evaluates location evidence against a location claim, producing a credibility assessment.

- **Credibility assessment** — the output of an evidence function: an estimate of where and when the entity was, given the evidence — together with qualifiers describing how robust that estimate is and under what assumptions it holds (see the analytical sketch below).

This framework distinguishes between location evidence (portable, objective artifacts) and credibility assessments (contextual evaluations). This separation acknowledges that what constitutes sufficient evidence depends on context — the same evidence might yield different assessments for different verifiers or applications. For example, the same location evidence may be highly valued by nation state A, but heavily discounted by nation state B, depending on the LVS used to produce it, and evidence supporting low-precision verification (i.e. a compute cluster) may be useless to one with higher precision requirements (delivery). By

1. "Verifiable" here means a downstream party can check the output's origin and integrity — and in some designs its freshness — without trusting the prover. Note that no LVS verifies a location by itself: a stamp is not a verified location; a decision is. Verification completes downstream, when a verifier evaluates evidence against a claim.
2. Signals may themselves be authenticated (signed navigation messages, trusted peripherals) or unauthenticated; a core design principle is to push cryptographic verifiability as far up the signal supply chain as possible — Galileo's OSNMA, which authenticates navigation messages at their source, is one example.

standardizing how evidence is structured while allowing flexibility in how it's evaluated, we can support diverse use cases from low-stakes applications to high-security scenarios.

The framework is also meant to support an end-to-end verifiable pipeline. The aspiration is the same as in a cryptographic protocol: a verifier receiving an evidenced claim should be able to derive confidence in specific properties — the origin and integrity of each stamp, the authenticity of the signals behind it, and so on — from the construction itself, using cryptographic and hardware techniques, with trust assumptions made explicit rather than implicit.

The Location Verification Lifecycle

In our framework, the verification lifecycle involves:

1. Some **location claim** is made by a prover.[1]
2. **Signals** produced by signal sources are processed by stamp generators into signed **location stamps**, wherever those generators run: on a device in the field, in third-party infrastructure such as witness or landmark networks, or in an inspector's toolkit.
3. Stamps are composed into **location evidence** and attached to the claim, forming a **location manifest**.
4. An **evidence function** evaluates the evidence against the claim, yielding a **credibility assessment**.
5. A **decision rule** maps the assessment to a decision.

Stamps need not touch the prover's infrastructure at all: an inspector who visits a site and independently files a report is generating location evidence like any other system.

For emphasis, this work operates as an abstraction layer *above* individual location verification systems — it complements and depends on them for generating location evidence. We're in the early stages of formalizing this framework and are seeking input on any and all of it.

3. Analytical Sketch

This section sketches how we might formalize this intuition mathematically, as a basis for measurement and simulation.

To generate credibility assessments, we apply an **evidence function** E to evaluate evidence against claims:

$$\mathrm{E} : (C, E) \mapsto A$$

The assessment A has two parts:

$$A = (\pi, Q)$$

where π is a **spatiotemporal posterior** — a probability distribution over where and when the event took place, given the evidence — and Q is a vector of **qualifiers** describing the robustness of that posterior: the independence of the contributing systems, the cost of forging the evidence behind it, its privacy properties, its degree of decentralization, and so on. The posterior says what the evidence indicates; the qualifiers say how far the indication can be trusted, and against which adversaries. The same location manifest can yield different assessments, depending on the evaluation function.

This provides a structured way to compare or combine heterogeneous sources of location evidence, regardless of how they are generated. The sections below build up the objects involved, from ground truth to decision.

1. Reality

We posit an underlying, true spatiotemporal mapping:

1. Note this is often implicit — the claim may be tacit in the interaction, or derived from the signals themselves.

$\ell_{\text{true}} : X \times \mathrm{T} \to \mathrm{S}$

where

- X is the set of entities or events,
- T is the set of times, and
- S is the spatial domain.

For any event $x \in X$ and time $t \in \mathrm{T}$, $\ell_{\text{true}}(x, t)$ denotes the true location of x at t.

The mapping is only locally observable; each observation yields bounded evidence by which claims about ℓ_{true} may be evaluated.

2. Claim

A prover asserts a claim about an event:

$C = (x, L, T)$

interpreted as the statement:

$\exists\, t \in T : \ell_{\text{true}}(x, t) \in L$

In words: event x occurred in region L at some time within interval T. This expresses what the prover asks the verifier to believe. Note that L and T are extents, not points: a claim always refers to an envelope of spacetime.

3. Signals

To compile evidence to corroborate C, the prover draws on one or more location verification systems. Each system's signal sources yield a set of raw observables concerning the entity's location.

$O_i = o_{i1}, o_{i2}, \ldots, o_{ik}$

Observables may include direct measurements (physical signals, environmental identifiers) or signed assertions from external entities (peers, authorities, services with location knowledge).[1]

4. Location stamps

A stamp generator function processes signals into a signed, verifiable artifact:

$s_i = g_i(O_i)$

where $g_i : O_i \to s_i$ encodes the system's localization model. Systems typically determine location through:

- **Inference**: computing position from physical signals (trilateration, timing analysis)
- **Reference**: looking up observed identifiers in location databases (IP geolocation, WiFi mapping)
- **Corroboration**: providing location-relevant attestations or patterns (peer witnesses, authority vouching, behavioral data)

Each stamp must be digitally signed to make its origin and integrity verifiable.

1. Observables vary in how much of their integrity can be checked: some arrive authenticated at the source (signed navigation messages, challenge–response exchanges), some are protected in transit or at the sensor interface (trusted peripherals), and some are bare measurements whose integrity rests entirely on the device that took them.

Beyond this minimal requirement, systems may include additional durability assurances — mechanisms that make falsification or replay more costly or detectable, such as cryptographic attestations, hardware roots of trust, or economic costs.

Most existing verification systems provide only partial guarantees; mapping and comparing these durability properties remains an active area of research.

5. Location evidence

A location evidence bundle is composed of m stamps:

$$E = s_1, s_2, \ldots, s_m$$

Each stamp carries its own provenance, and may carry uncertainty estimates and explicit trust assumptions. Together, these stamps form the evidentiary basis for a specific spatiotemporal claim. The location evidence and location claim are packaged as a location manifest.

6. Evidence function

An evidence function E evaluates location evidence E against a claim C to produce a credibility assessment:

$$\mathrm{E} : (C, E) \mapsto A = (\pi, Q)$$

The posterior π is a probability distribution over spacetime $\mathrm{S} \times \mathrm{T}$, conditional on the evidence and on a threat model θ — the assumptions the verifier makes about what an adversary can do (spoof unauthenticated signals, collude with peers, compromise a signing key). The credibility of the claim itself is the posterior mass inside the claim's envelope:

$$\Pr[C \mid E, \theta] = \pi(L \times T)$$

— the probability, given the evidence and under the assumed threat model, that the event occurred within L during T.[1] The qualifiers Q record information this number compresses away: which threat models the evidence is resistant to, at what forgery cost, with what privacy and decentralization properties.

In practice, E considers multiple factors when forming π and Q:

- $\mathrm{correlation}(E)$ — *independence:* the degree to which constituent systems provide distinct information rather than redundant signals.
- $\mathrm{strength}(E)$ — *robustness:* the intrinsic reliability of each contributing system, including calibration accuracy, error bounds, and resistance to forgery or manipulation.
- $\mathrm{relevance}(C, E)$ — *fit to the claim:* how directly the evidence bears on the claim's entity, region, and interval. A stamp about the right device but the wrong day has zero relevance; a stamp localizing to city scale has only partial relevance to a building-scale claim.

How π and Q should be computed is unsettled, and we are not committed to a formalism. Bayesian updating is the obvious starting point for combining independent measurements. Beyond it, several bodies of theory appear relevant — Dempster–Shafer theory (Shafer, 1976) and subjective logic (Jøsang, 2016) both extend probabilistic reasoning to handle conflicting evidence, ignorance, and provenance-dependent trust — but we have not evaluated their fit to composable location evidence, and flag them here as candidates rather than commitments (see open question 1).

---

1. This treats x as a point event — something that happened once, somewhere. Claims about persistent entities differ in kind, not just duration: "the chip remained within L throughout T" concerns a path through spacetime rather than a point, and the natural object becomes a distribution over trajectories. We defer this to future work (see open question 1).

Evidence functions form a family — each verifier selects or constructs its own. Characterizing that family is itself an open problem: what properties must any admissible E satisfy (the design properties below are a first attempt), and how are particular members built, validated, and compared? A key challenge is to characterize what the qualifiers in Q should capture — accuracy, integrity, cost of forgery, privacy, decentralization, and so on — and to establish methods for quantifying these across heterogeneous location verification systems with different assurance models and interdependencies.

Application-specific evidence functions

Different applications require fundamentally different evidence functions. The evidence function is selected by the verifier, and some verifiers may exclude entire categories of evidence based on their trust models:

- **Compute-governance settings** might require hardware-rooted attestations and reject self-reported measurements entirely.
- **Regulatory-compliant systems** might require government attestations and reject peer-to-peer evidence.
- **Privacy-preserving applications** might only accept evidence with specific privacy guarantees, rejecting evidence that exposes raw location data.
- **Decentralized infrastructure networks** (such as blockchain validator sets) might reject any evidence from centralized authorities, however difficult it is to forge.

Depending on a verifier's requirements, some evidence types may be excluded altogether. For example, a location stamp derived from a proprietary indoor-positioning network — such as one operated by a major mapping or retail-analytics company — might deliver sub-meter accuracy and be difficult to forge without privileged access, yet be assigned zero credibility by systems that reject centralized data services. Likewise, a state verifier in a compute-governance setting might assign zero weight to latency measurements taken from landmark nodes operated by a geopolitical adversary.

Evidence functions thus encode not just how to evaluate evidence, but which evidence to consider. This application-specific filtering reflects the fundamental trust assumptions of each use case.

7. Decision rule

To make credibility assessments actionable, verifiers need to move from the assessment $A = (\pi, Q)$ to a decision.

A decision rule combines the claim credibility $\Pr[C \mid E, \theta]$ with the qualifiers Q, typically by weighting them into a value that can be compared against an application-specific threshold:

$$w : (\Pr[C \mid E, \theta],\ Q) \mapsto p \in [0, 1]$$

Different applications weight the components differently:

- **High-security applications** might require $Q_{\text{forgery-cost}} > \theta_{\text{high}}$ regardless of spatial precision
- **Navigation services** might weight the precision of the posterior — how narrowly it localizes the event — heavily while tolerating lower forgery resistance
- **Decentralized systems** might require $Q_{\text{decentralization}} > \theta_{\text{min}}$ as a hard constraint before considering other dimensions

The decision rule encodes these priorities, transforming the rich credibility assessment into a context-specific decision value.

Importantly, the framework doesn't prescribe a universal decision rule. What constitutes "sufficient evidence" is inherently application- and verifier-specific. The same credibility assessment might be accepted by one verifier and rejected by another, based on their different requirements and risk tolerances.

This separation between evidence evaluation (producing A) and decision-making (applying w) allows the framework to serve diverse applications while maintaining a common language for location verification.

4. Design Properties

This framework has several properties that make it useful across different contexts.

Evidence responsiveness

Every valid, independent signal should update the assessment — no admissible evidence is ignorable. This cuts both ways: corroborating evidence should strengthen support for a claim, and conflicting evidence should widen uncertainty rather than be suppressed. In general, accumulating independent evidence should sharpen the assessment, tightening our understanding of whether the claim is truthful.

Redundancy discounting

When two signals convey the same information, the second does little to increase **certainty**. The framework should recognize correlation so that only genuinely independent evidence contributes to a more reliable assessment.

Forgery-cost condition

Location evidence is only as good as the cost of faking it. For any given application, the expected cost of forging all contributing signals should exceed the value of whatever action or transaction the verification underwrites. This keeps incentives aligned and makes deception uneconomical.[1]

Composability

Different location stamps can be combined into evidence bundles in a consistent way. The framework ensures that evidence from multiple sources can be meaningfully integrated, with each stamp's contribution properly accounted for based on its independence and relevance. This makes it possible to layer multiple verification systems together while maintaining interpretability.

Verifiability

As introduced above, the framework is designed to support verifiable computation throughout the lifecycle. Location stamps should be cryptographically signed, evidence bundles should maintain clear provenance, and evaluation functions should be identifiable so their outputs can be independently reproduced or verified. In practice, systems will vary in how far they go — from fully transparent, publicly reproducible evaluation to zero-knowledge proofs or trusted execution — but all should make explicit what can be verified and what trust assumptions remain.

Neutrality

The framework is unopinionated about how location evidence is generated, how assessments are evaluated, or how evaluations are used. It accommodates diverse location verification systems and evidence functions, and it makes no assumptions about which attributes verifiers value most. It supports diverse modes of decision-making — from simple thresholds to adaptive or collective reasoning — allowing applications to define what "sufficient evidence" means within their own operational logic. Different implementations can compete, benchmark, and evolve within a common architecture.

---

1. This condition conceals a game-theoretic difficulty: the verification threshold *should* scale with the prover's payoff from lying, but that payoff is the prover's private information. The verifier is provisioning security against a gain it cannot observe and may underestimate. We return to this in the open questions.

Together, these properties describe how evidence should behave as it accumulates: more independent, high-integrity signals sharpen our picture of reality, while noise, redundancy, or manipulation are naturally constrained.

5. Illustrative Examples

Static deployments: a chip in a data center

Consider an AI accelerator whose operator asserts an honest location claim $C = (x, L, T)$: the chip is operating within jurisdiction L throughout interval T. Evidence is collected from three largely independent location verification systems:

- a **network-measurement system**, which bounds the chip's possible location through latency measurements from geographically distributed landmark nodes — round-trip times constrain distance under speed-of-light and network-propagation limits (Brass & Aarne, 2024). Crucially, each challenge–response is signed by the chip's on-die hardware root of trust, binding the measurements to that specific device;
- an **authenticated-GNSS system**: a receiver in a secure module producing fixes from authenticated navigation signals (e.g., Galileo OSNMA). Signal authentication raises the cost of spoofing substantially (Fernández-Hernández et al., 2016), though record-and-replay attacks remain feasible for sophisticated adversaries (Wang et al., 2025 — which is precisely why it serves as one stamp among several rather than a guarantee on its own; and
- **periodic physical-inspection attestations** (as in export-control end-use checks) add a third, institutionally anchored stamp.

Each system produces a signed **location stamp**. Combined, they form a location manifest (C, E) with $E = s_{\text{network}}, s_{\text{gnss}}, s_{\text{inspector}}$.

Applying an evidence function yields a credibility assessment $A = \mathrm{E}(C, E)$ reflecting:

- **Low correlation**: latency measurement, authenticated satellite positioning, and institutional inspection rest on largely independent mechanisms, and sit in three different trust domains: verifier-side network infrastructure, the satellite system plus a secure module, and a human institution,
- **High robustness**: all three are costly to forge,
- **High relevance**: all three directly measure the claimed device during the claimed interval.

The expected cost of forgery rises sharply: an attacker must simultaneously defeat latency measurements from many independent vantage points *and* mount a replay attack against authenticated satellite signals *and* subvert the hardware and human elements of the institutional inspection regime — a significantly higher bar than attacking any of the systems alone.

Mobile deployments: a delivery drone

Now consider a claim with much higher precision requirements: a drone operator asserts that a package was delivered at a specific address L within interval T. Latency-based techniques, useful at metropolitan or national scale, cannot resolve location at this precision. The drone instead composes evidence across several complementary systems:

- an **onboard navigation stamp**: the drone's fused GNSS, inertial, and barometric solution, cross-checked internally for spoofing consistency;[1]
- an **RF-fingerprinting stamp**: the observed radio environment (WiFi and cellular identifiers, signal strengths) matched against reference maps of the delivery area;
- a **computer-vision stamp**: a location-extraction model matching delivery imagery against geo-referenced views of the claimed address (e.g., StreetCLIP; Haas et al., 2023);

1. Sensor fusion is itself compositional, but composition occurs within a single trust domain. However many sensors feed it, the fused solution is one stamp because it is signed in one on-device secure hardware component.

- a **recipient attestation**: the customer confirms receipt by providing a PIN at the delivery point — a social corroboration of the physical event.

Individually, several of these stamps are weaker than the data-center evidence above: GNSS can be spoofed, RF maps are constantly changing, recipients can collude. The first three are also generated onboard and signed by the same device — they share a root of trust, and against a compromised drone they fail together. Here, composition offers some protection against an over-the-air adversary: spoofing satellite signals, fabricating a consistent radio environment, and forging imagery that matches the claimed scene, all coherently for the same place and time, is expensive. Which adversaries a composition resists is what the qualifiers Q record. The recipient attestation — the only stamp generated outside the drone's trust domain — matters more than its individual strength suggests.

Both cases illustrate the framework's key insight: independent evidence sources compose into stronger assessments, while redundant sources add little value.

6. Open Questions

This framework raises more questions than it settles. The ones we consider most pressing:

1. **Structure of the credibility assessment:** Is the two-part formulation — a spatiotemporal posterior plus robustness qualifiers — the right shape? Should the posterior be computed per threat model, yielding a family of distributions rather than one? How should claims about persistent entities — paths through spacetime, not point events — be handled? And which formal machinery fits best: Bayesian updating, Dempster–Shafer belief functions, subjective logic, or some combination? This is where we most expect the framework to evolve.

2. **Independence and correlation:** How do we measure and model correlation between verification systems, so that composite assessments meaningfully reflect independence rather than redundancy?

3. **Adversarial robustness:** How robust can location verification be when the prover itself is motivated to deceive — a data-center operator misreporting where chips run, or an autonomous system spoofing its own sensors? What assumptions break down under these threat models, and what mechanisms could reestablish confidence? It may help to index threat models to established attacker-capability scales — such as the OC1–OC5 operational capacity categories used in work on securing model weights (Nevo et al., 2024) — and ask which evidence compositions survive each class.

4. **Durability mechanisms:** What cryptographic, hardware, and economic techniques can increase the durability of location evidence — making falsification or replay detectably costly — and how can these be compared or composed across systems?

5. **Calibrating to adversary incentives:** The forgery-cost condition says verification strength should exceed the prover's payoff from lying — but that payoff is private information. How should verifiers provision security against a gain they cannot observe? What mechanisms (disclosure, bonding, tiered verification) could surface or bound it?

6. **Benchmarking strength:** How can we compare and benchmark the strength of different location verification systems — including their accuracy, durability, and cost of forgery — in a consistent, reproducible way?

7. **Privacy and disclosure:** What privacy affordances are essential across use cases, and which privacy-preserving techniques remain compatible with verifiable evaluation?

8. **Governance integration:** If composable location evidence becomes available, how might governance regimes use it — for compute oversight, operational restrictions on autonomous systems, or compliance auditing — without introducing new surveillance or centralization risks?

9. **Where to begin:** Which existing verification systems or datasets are ready for early experimentation, and what minimal testbeds or metrics would yield the clearest signal on feasibility and value?

7. Next Steps

If you're working on hardware security, compute governance, autonomous-systems assurance, network measurement, or simulation frameworks that could integrate spatial evidence, we'd like to compare notes. Our next step is formalizing testable evidence functions and benchmarking existing location verification systems.

References

Brass, A., & Aarne, O. (2024). *Location verification for AI chips*. Institute for AI Policy and Strategy. https://www.iaps.ai/research/location-verification-for-ai-chips

Chip Security Act, H.R. 3447, 119th Cong. (2025).

Fernández-Hernández, I., Rijmen, V., Seco-Granados, G., Simón, J., Rodríguez, I., & Calle, J. D. (2016). A navigation message authentication proposal for the Galileo open service. *NAVIGATION: Journal of the Institute of Navigation, 63*(1), 85–102.

Haas, L., Alberti, S., & Skreta, M. (2023). *Learning generalized zero-shot learners for open-domain image geolocalization*. arXiv:2302.00275.

Jøsang, A. (2016). *Subjective logic: A formalism for reasoning under uncertainty*. Springer.

Kabatnik, M., & Zugenmaier, A. (2001). Location stamps for digital signatures: A new service for mobile telephone networks. In *Networking — ICN 2001* (LNCS 2094). Springer.

Nevo, S., Lahav, D., Karpur, A., Bar-On, Y., Bradley, H. A., & Alstott, J. (2024). *Securing AI model weights: Preventing theft and misuse of frontier models*. RAND Corporation, RR-A2849-1.

Shafer, G. (1976). *A mathematical theory of evidence*. Princeton University Press.

Wang, H., Zhang, Y., Zhu, X., He, J., Zhao, S., Shen, Y., & Jiang, X. (2025). *Practical spoofing attacks on Galileo Open Service Navigation Message Authentication*. arXiv:2501.09246.

Many thanks to Adam Spiers, Taylor Oshan, Andy Tudhope, Vivek Singh, Ron Erlih, Seth Docherty, Kiersten Jowett, Holly Grimm, j-mars, jason-james, and Danny Gattman for all their contributions to this work.