

ವಿಶ್ವೇಶ್ವರಯ್ಯ ತಾಂತ್ರಿಕ ವಿಶ್ವವಿದ್ಯಾಲಯ, ಬೆಳಗಾವಿ

VISVESVARAYA TECHNOLOGICAL UNIVERSITY, BELAGAVI



BANGALORE INSTITUTE OF TECHNOLOGY

DEPARTMENT OF INFORMATION SCIENCE AND ENGINEERING

K.R.ROAD, V.V.PURAM, BANGALORE - 560004

Deep Learning Project

GIST: Weakly-Supervised Detection, Classification and Localization of GNSS Jamming and Spoofing from Crowdsourced ADS-B

Presented by:

NAME	USN
Vansh Agarwal	1BI23IS133
Vichitra Attri	1BI23IS135
Vidya Bai K	1BI23IS137

Under the guidance of:

Dr. Asha T.

Deep Learning course

GIST

Every aircraft as a remote GNSS sensor: detecting, classifying and localizing GPS interference from crowdsourced ADS-B, with no manual labels

Transformer · Graph attention · Weak supervision · Differentiable physics

OUTLINE

01 Abstract and motivation

02 Literature survey and research gap

03 Problem statement and objectives

04 Dataset and licensing

05 System architecture

06 Physics features and weak supervision

07 GIST model and localization head

08 Simulator and experimental setup

09 Results, ablations and case study

10 Challenges, limitations, conclusion and future work

Abstract

GIST is a deep-learning system that detects GNSS (GPS) interference affecting civil aircraft, tells jamming apart from spoofing, and estimates where a jammer is. It uses only what aircraft already broadcast over ADS-B, collected by volunteer receivers and published as open data (adsb.lol, ODbL).

Public maps threshold one integrity field (NIC) per hexagon per day. That catches jamming but misses spoofing, because a spoofed receiver reports high integrity. GIST instead computes physics-consistency features a spoofer cannot fake cheaply (barometric vs GNSS altitude against neighbours, wind consensus, position jumps, co-located aircraft at different altitudes), encodes each aircraft's 10-minute track with a Transformer, shares evidence between nearby aircraft with spatial graph attention, and localizes jammers with a differentiable Friis / radio-horizon likelihood.

No labels are drawn by hand: ten physics-based labeling functions are combined by a Dawid–Skene label model, and synthetic jammers and spoofers injected into real traffic provide exact targets.



Detects

Interference per aircraft and 10-minute window, minutes after it starts.



Classifies

Normal, jammed or spoofed, including spoofing that NIC thresholds cannot see.



Localizes

Jammer position from which aircraft are affected, through a differentiable physics layer.

Jamming vs Spoofing: Why NIC Maps Miss Half the Problem

Jamming

Signal denied by RF noise

- Receiver loses lock; broadcast NIC collapses toward 0
- Position reports drop out while altitude and track continue
- Visible in integrity fields, so NIC-threshold maps (gpsjam.org) catch it
- Location of the emitter is still unknown

Spoofing

Counterfeit signal, confident lie

- Receiver locks onto fake satellites and reports high integrity (NIC ≥ 8)
- Positions jump, circle or drift; GNSS altitude disagrees with barometric altitude
- Integrity-threshold maps are structurally blind to it
- More dangerous: the crew may trust the wrong position

Why now. Interference affecting aviation has grown sharply since 2022 around the Baltic, Eastern Mediterranean, Black Sea, Middle East and India–Pakistan border. In our first day of data, the share of points with NIC ≤ 6 was 15.8 % over the Baltic versus 1.7 % over SE Australia. Barometric altitude and air data do not depend on GNSS, and a spoofer cannot make hundreds of aircraft mutually consistent: that redundancy is what GIST learns from.

Literature Survey

#	Work	Approach	Gap for our problem
[1]	Liu, Lo, Walter — ION ITM 2022	ADS-B jamming simulator; Friis model; tolerable power ≈ −115 dBW	Jamming only; needs clean affected / unaffected labels
[2]	Liu et al. — NAVIGATION 72(3), 2025	Non-linear least squares on affected aircraft; ~4 km on Denver 2022	No spoofing; sensitive to benign low-NIC aircraft (we measure this)
[3]	Schäfer et al. — Jammer on the Horizon	Radio-horizon intersection from ADS-B coverage	Localization only; assumes jamming is already detected
[4]	Wiseman — gpsjam.org	Daily H3 hexagon map, share of aircraft with low NIC	Day granularity; cannot see spoofing; no emitter location
[5]	Eichstaedt et al. (DLR) — IEEE SHaRC 2026	Global interference measurement from space-based ADS-B	Measurement, not learned classification or localization
[6]	Aireon — Spoofing white paper, 2026	NIC and position cross-checks on space-based ADS-B	Industry rules; no learning across aircraft
[7,8]	Ratner et al. 2017; Dawid & Skene 1979	Data programming; EM label model over noisy labelers	Method we adopt to build labels without annotation

Re-checked against 2025–2026 literature in Phase 1 of the project (LOG.md). Full references on the References slide.

Disadvantages of Current Systems

No prior work jointly separates jamming from spoofing with physics invariants, learns across aircraft, localizes with a learned + physics hybrid, and needs no manual labels.

Blind to spoofing

Integrity-threshold maps flag low NIC; spoofed aircraft report high NIC.

Day-level latency

Public maps aggregate per day; an event is visible only afterwards.

No emitter location

Detection maps say where aircraft were affected, not where the source is.

Labels are scarce

Ground-truth events with known emitter positions are rare and not public.

Aircraft treated alone

Per-aircraft rules ignore that neighbours see the same emitter.

Fragile localizers

Least-squares needs clean affected sets; benign low-NIC aircraft corrupt it.

Problem Statement

> *Given only crowdsourced ADS-B broadcasts and no manual labels, detect GNSS interference per aircraft within minutes, classify it as jamming or spoofing, and estimate the jammer's location, robustly to aircraft whose low integrity has nothing to do with the emitter.*

Inputs and outputs

Input
Per-aircraft ADS-B traces: position, baro and GNSS altitude, speed, track, NIC / NACp / SIL, readsb wind

→

Scene
All aircraft in a ±3° latitude box over one 10-minute window (≤ 256 aircraft, 20 steps of 30 s)

→

Output
Class per aircraft-window and per step (normal / jammed / spoofed) + jammer location

1. **No ground truth**
   labels must come from physics and simulation

2. **Rare spoofing**
   few spoofed windows per day

3. **8 GB laptop**
   Apple M3; every model sized to fit

Objectives

01

Build an openly licensed, weakly-supervised dataset of jamming / spoofing labels from physics-based labeling functions and a label model.

02

Design spoofing-sensitive features that do not rely on the integrity fields a spoofer controls.

03

Train GIST: a trajectory Transformer with spatial graph attention for 3-way classification per aircraft and per step.

04

Localize jammers with a differentiable Friis / radio-horizon head trained sim-to-real on synthetic emitters in real traffic.

05

Compare fairly against re-implemented baselines (gpsjam threshold, Isolation Forest, LSTM-AE, NLS, radio horizon) on day-split data.

06

Report honestly: 3 seeds, ablations, independent MLAT check, and the limits of weak labels and simulation.

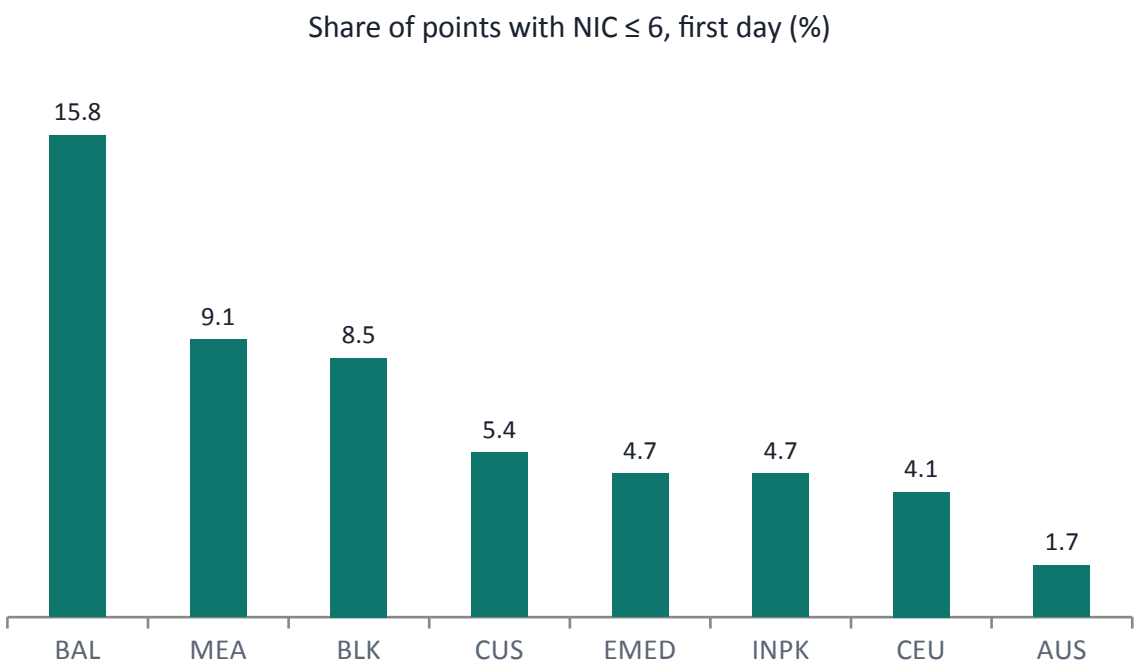
Dataset: Open, Crowdsourced ADS-B (adsb.lol, ODbL)

67 M
airborne trace points
in study regions (interim, 5 days)

1.22 M
aircraft-windows
10 min, 20 × 30 s steps

67,333
aircraft in one day
global, 2026-10-25

114 s
to parse one day
6 workers, 340 MB RAM

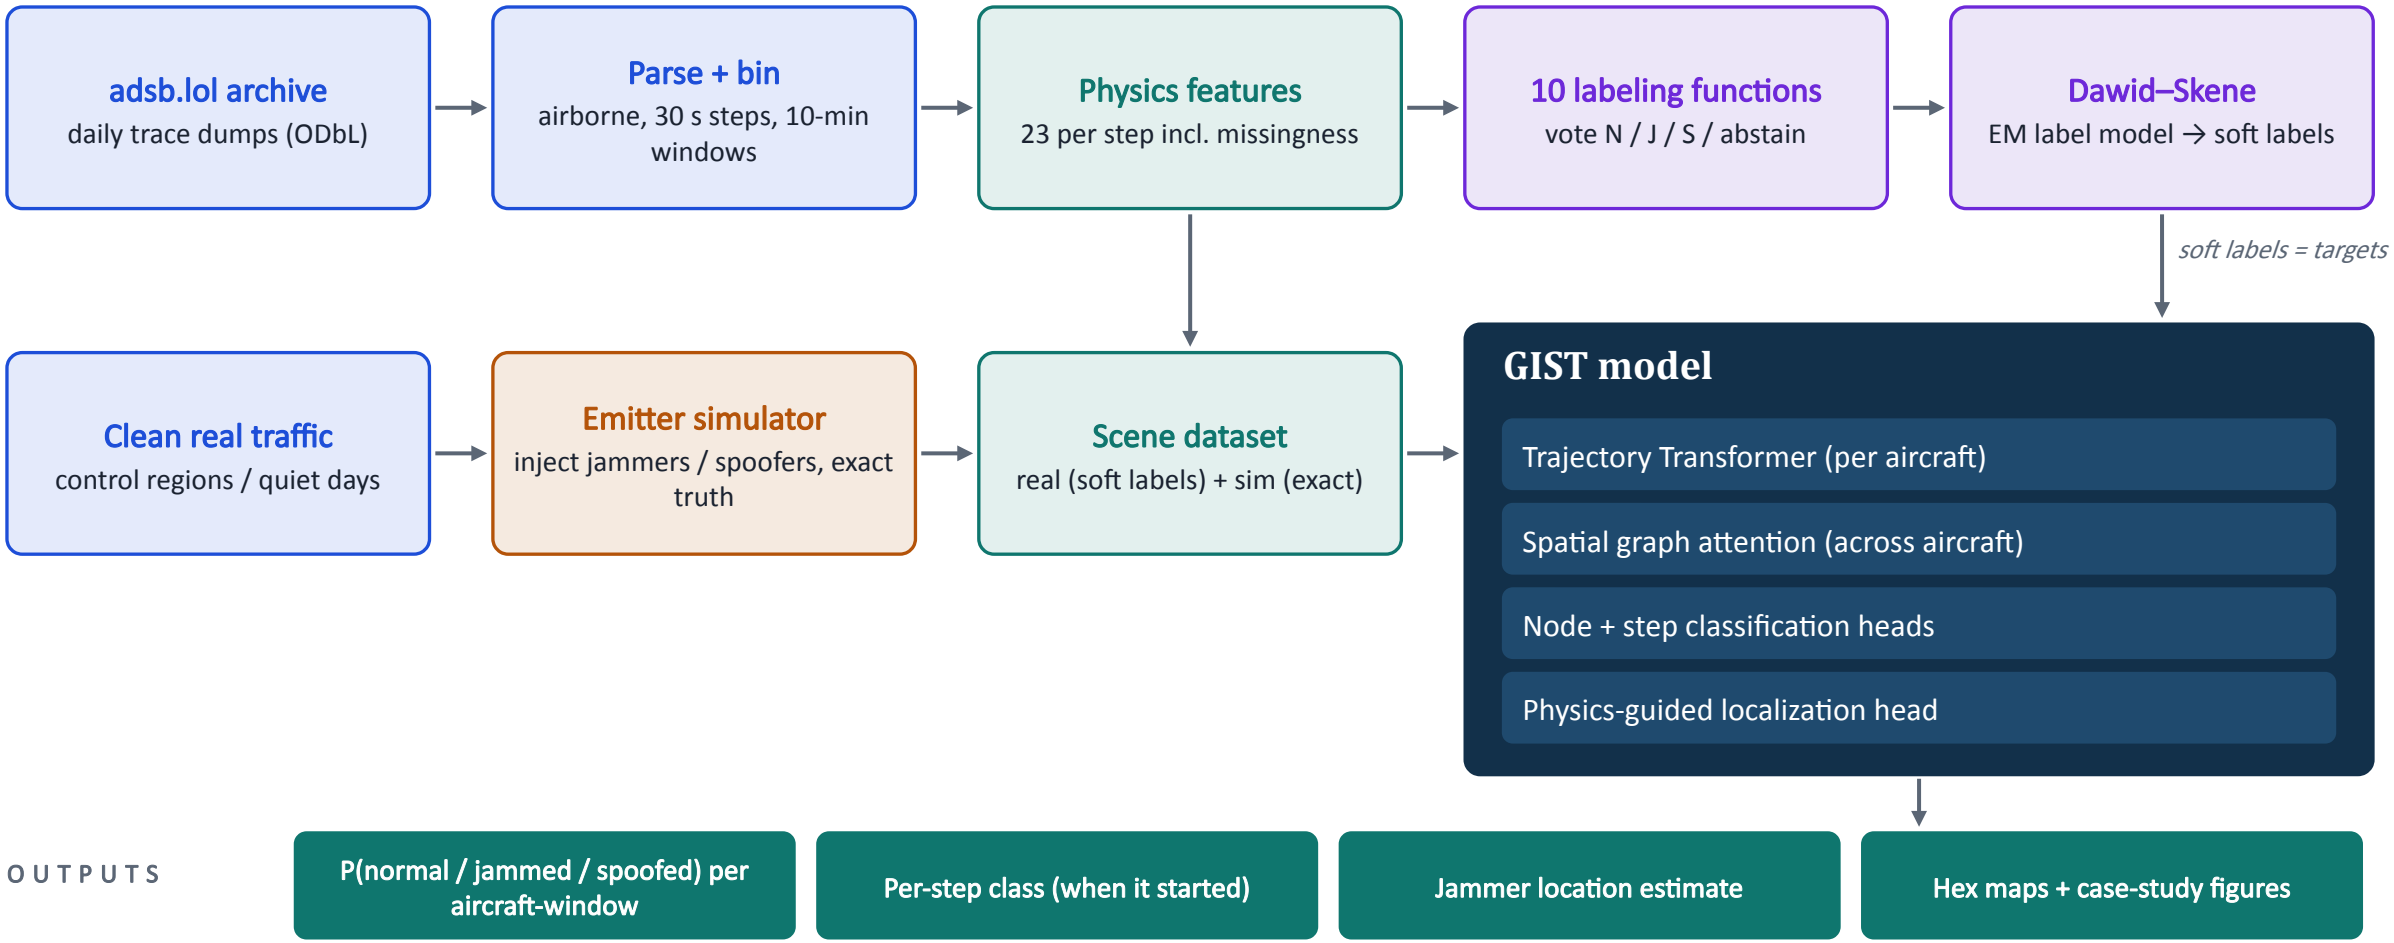


Source and licence
adsb.lol globe_history: daily readsb trace dumps, ODbL-1.0; attribution "ADSB.lol and contributors". ADS-B Exchange samples rejected (evaluation-only terms bar derivative works).

Regions
Interference: BAL (Baltic), EMED, BLK (Black Sea), MEA (Iraq / Iran / Gulf), INPK (India–Pakistan), MYN. Comparison: CEU (low-interference Europe), clean controls CUS (US Midwest) and AUS (SE Australia).

Days
1st of each month, Jan 2024 → Oct 2026, plus 2026-10-25; the final pipeline is still downloading them. Point mix: ADS-B 95 %, ADS-R 2 %, MLAT 2 %, TIS-B < 1 %.

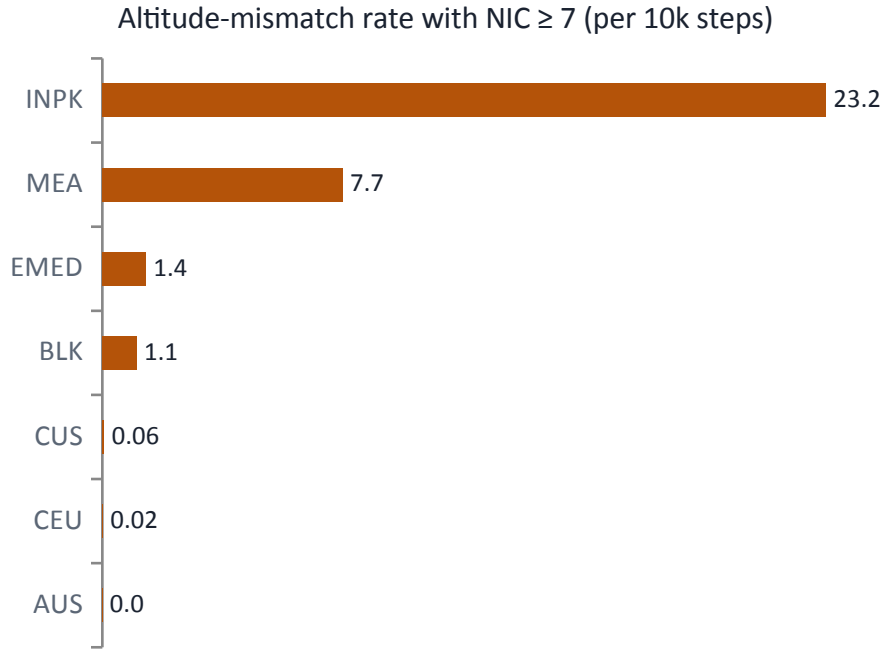
System Architecture



Physics-Consistency Features

Feature (per 30 s step)	What it catches	Can a spoofer fake it?
NIC / NACp / SIL / GVA and their drops	integrity collapse = jamming	Yes
Kinematic residual $\lvert \lVert \Delta p \rVert - v \cdot \Delta t \rvert$	position jumps	Partly
GNSS − baro altitude vs neighbour median	GNSS altitude drift	No (baro is independent)
readsb wind vs neighbour consensus	broken wind triangle	No
Spoof cluster: ≥ 3 aircraft in one 0.005° cell, ≥ 5000 ft apart	"spoof point" sharing	No
ADS-B gaps, stale flags, position loss	jamming drop-outs	—

Neighbour consensus = trusted aircraft (NIC ≥ 7) in the same H3 res-3 cell, 5-min slot and 6000 ft band.

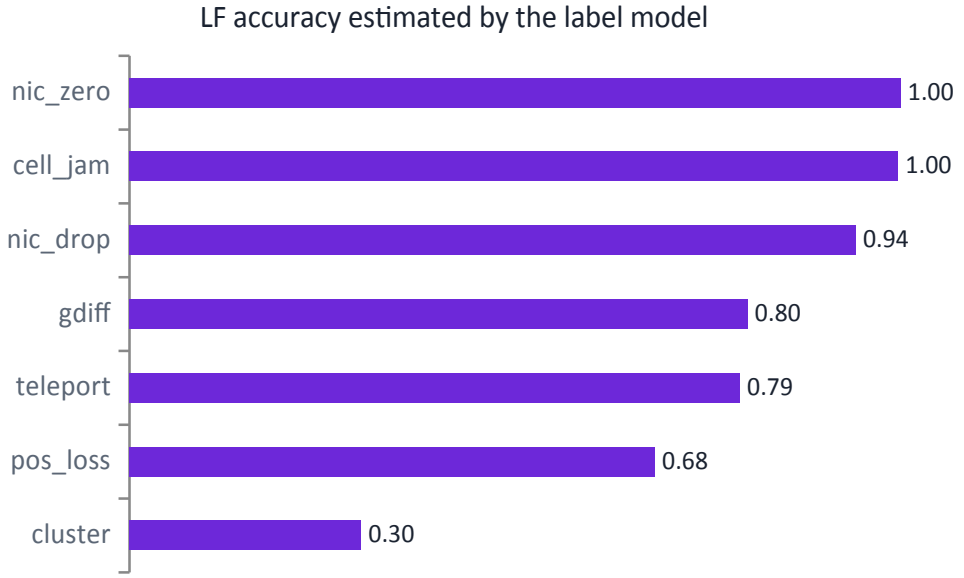


Validated before use. Spoof-sensitive features fire in spoofing regions while integrity looks fine, and almost never in controls. Naive wind from tas/heading was useless (p99 of several hundred kt) and is kept as a model feature only, not a labeling function.

Weak Supervision: Labeling Functions + Dawid–Skene

LF	Vote	Condition (10-minute window)
nic_zero	JAM	modern, cruise, ≥ 30 % steps NIC ≤ 1
nic_drop	JAM	cruise, NIC ≥ 8 seen and ≥ 30 % steps NIC ≤ 6
pos_loss	JAM	cruise, ADS-B gap ≥ 90 s, ≥ 10 % steps NIC ≤ 6
cell_jam	JAM	≥ 30 % of cell aircraft and ≥ 30 % own steps NIC ≤ 6
teleport	SPOOF	position jump > 10 km with NIC ≥ 7
gdiff	SPOOF	≥ 2 steps |Δh geo−baro − nbr| > 1000 ft, NIC ≥ 7
cluster	SPOOF	member of co-located, altitude-spread cluster
mlat †	SPOOF	MLAT vs ADS-B disagreement > 10 km
clean	NORMAL	all NIC ≥ 8, no jump / gdiff / cluster, ≥ 10 steps
clean_lowcell	NORMAL	modern, cell NIC-low share < 5 %, own NIC never low

† held out of training labels; evaluation only. Region names are never used inside LFs.



- EM estimates each LF's confusion matrix per class; output = soft labels for soft cross-entropy.
- Coverage 77–92 % of windows in every region; controls CUS 99.9 % / AUS 99.6 % normal.
- Bug found and fixed: all-abstain windows inherited a “spoofed” posterior (12.8 % in AUS) → now left unlabelled.

GIST Model Architecture

PER AIRCRAFT

20 steps × 23 features
30 s bins, missingness flags

↓

MLP embed + [CLS] + pos. emb.

↓

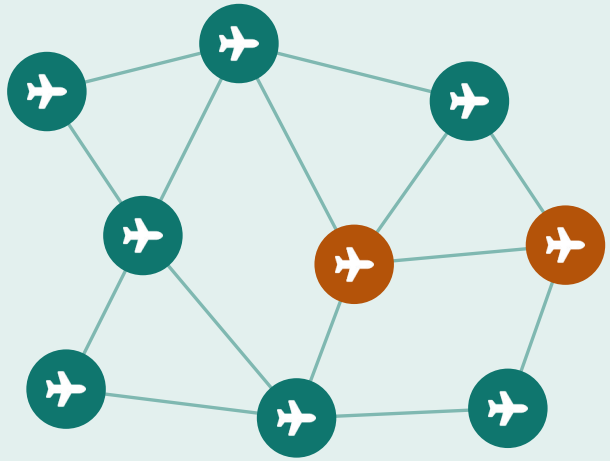
Transformer encoder
3 layers, pre-norm, d = 96, 4 heads

↓

node token + 20 step tokens

→

ACROSS AIRCRAFT (SCENE)



Spatial graph attention × 3
Aircraft within 400 km attend to each other (≤ 256 per scene). Logits get a learned bias from edge features (Δx, Δy, d, Δh), GATv2-style, in plain PyTorch.

HEADS

Node head
window class: normal / jammed / spoofed (soft cross-entropy on label-model targets)

Step head
per-step class from step token ⊕ node embedding (exact sim labels)

Localization head
physics likelihood + learned heatmap → soft-argmax (next slide)

Training: AdamW, lr 3e-4, batch 4 scenes, 15 epochs (final runs: 8), fp32 on Apple MPS. Loss = soft CE (node) + CE (step) + 0.5 × (Gaussian-target CE + L1) for location. Checkpoint 2.0 MB.

Differentiable Physics-Guided Localization

1

Step jam probabilities

p_it from the step head at aircraft positions x_it and heights h_it

2

Physics likelihood

for every grid cell g (21 × 21, 40 km) and EIRP e: received power P_r by Friis at L1, antenna gain vs elevation, 4/3-Earth horizon

3

Score

$\ell(g,e) = \Sigma\, p \log \sigma(z) + (1-p) \log \sigma(-z),\ \ z = (P_r - \tau) / T,\ \ \tau = -115$ dBW + learned offset

4

Fuse + soft-argmax

logsumexp over e, plus a learned heatmap from grid-to-aircraft cross-attention; location = soft-argmax

5

Refine

physics-guided coarse-to-fine search around the coarse estimate

Effect of refinement (val, seed 0)

25.1 km
coarse soft-argmax median

18.4 km
after physics refinement

Why a physics layer?

- Only simulated emitters have known positions; physics keeps the head sensible on real traffic.
- Uses soft probabilities, so uncertain aircraft count less than with a hard affected set.
- Memory: 41 × 41 grid needed 2.3 GB autograd; 21 × 21 + refinement needs 0.36 GB.

Emitter Simulator: Exact Targets in Real Traffic

Jammer

60 % of simulated scenes · EIRP 20–50 dBW

- Ground (2–50 m) or airborne (1–10 km)
- Friis + 4/3-Earth line of sight + elevation-dependent antenna gain
- Receiver tolerance $N(-115, 3^2)$ dBW per aircraft
- Above tolerance: NIC → 0, GNSS altitude lost, positions drop with prob. σ(m/3)
- Within 6 dB below: NIC in 1–6

Spoofer

40 % of simulated scenes · EIRP 10–40 dBW

- Capture when margin > −5 dB; capture is sticky
- Fake track: static point, circle or drift
- NIC forced ≥ 8 (realistic), GNSS altitude offset
- Ground speed / track / wind made consistent with the fake track
- Neighbour and cluster features recomputed after injection

Calibration. A first 1–100 W prior affected a median 0.4 % of aircraft per scene, which is unrealistic; with the final priors a median 41 % are affected (q10 5 %, q90 82 %). Sanity gate: our NLS re-implementation reaches 2.4 km median on noise-free ground jammers (11 of 12 under 6 km).

Experimental Setup

DAY-BASED TEMPORAL SPLITS (NO DAY IN TWO SPLITS)

Train	Validation	Test
Jan 2024 → Jun 2025	Jul → Dec 2025	Jan → Oct 2026

Baselines (re-implemented)

- gpsjam-style cell threshold (NIC share)
- Isolation Forest on window features
- LSTM autoencoder (reconstruction error)
- Localization: centroid, radio horizon, NLS (Liu et al.)

Metrics

- Macro-F1 over 3 classes
- Average precision for jam and spoof
- Spoofing recall at 1 % false-positive rate
- Localization median and p95 error (km)

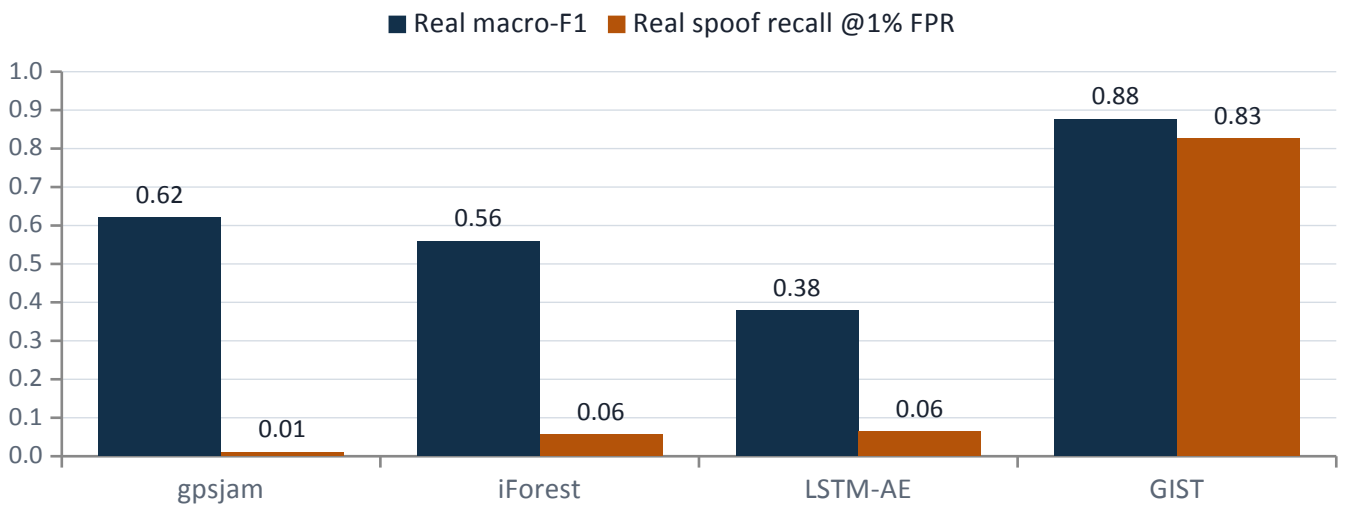
Compute

- Apple M3, 8 GB RAM, PyTorch on MPS
- MPS allocator watermarks + cache clearing
- Final: 3 seeds + 3 ablations, unattended
- ≤ 256 aircraft per scene, 21 × 21 grid

INTERIM · SEED 0 · TEST

Results: Detection and Jam-vs-Spoof Classification

Method	Real F1	Real Jam AP	Real Spoof AP	Real Spoof R@1%	Sim F1	Sim Spoof R@1%
gpsjam-style threshold	0.62	0.86	0.01	0.01	0.26	0.01
Isolation Forest	0.56	0.80	0.04	0.06	0.27	0.01
LSTM autoencoder	0.38	0.37	0.03	0.06	0.27	0.03
GIST *	**0.88**	**0.99**	**0.69**	**0.83**	**1.00**	**1.00**



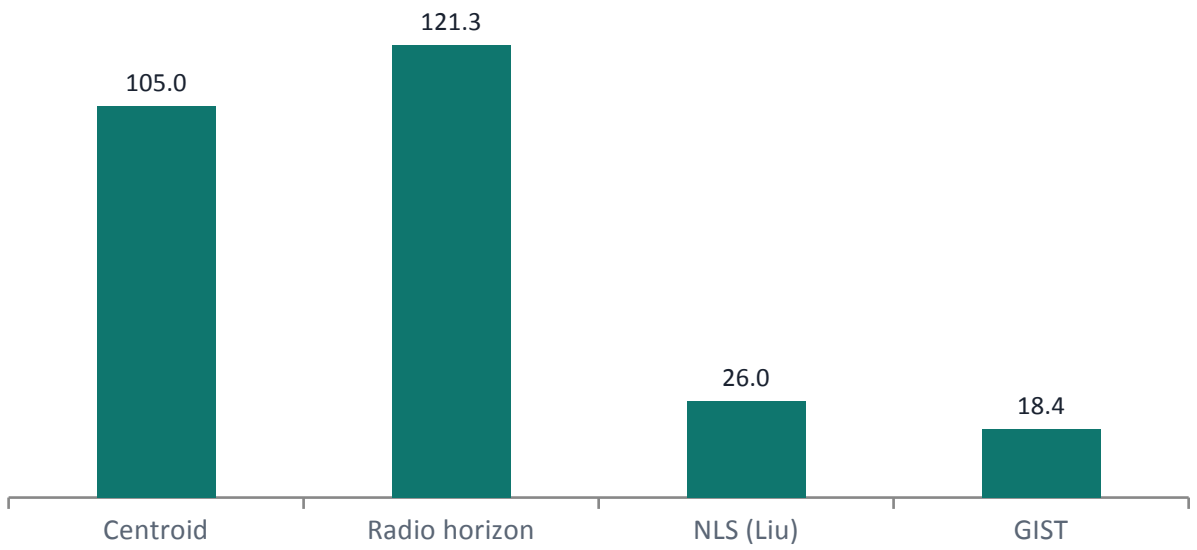
Reading the table

- The integrity threshold finds jamming (AP 0.86) but spoof recall is 0.01: it cannot see spoofing.
- GIST separates the two classes; spoof recall @1% FPR 0.83 on real traffic.
- * GIST = interim seed-0 model on the same 2026 test days as the baselines (n = 38,762 windows, 374 spoofed). Final 3-seed numbers will replace this row.

INTERIM · SEED 0 · TEST

Results: Jammer Localization and Distractor Robustness

Median error on simulated jammers, clean scenes (km, lower is better)



Method [km]	Clean median	Clean p95	+5 % benign low-NIC median	p95
Centroid of affected	105.0	260.2	126.4	266.4
Radio-horizon score	121.3	382.3	138.0	448.6
NLS (Liu et al.)	26.0	301.7	42.5	315.4
GIST *	**18.4**	**163.5**	**27.2**	**182.5**

** GIST interim: seed 0, on the same 2026 test scenes as the baselines (n = 334 clean, 160 distractor).*

Information limit. Even oracle step labels reach only 17.2 km median on the dev sims: receiver tolerance spread and unknown emitter height bound accuracy.

Distractor robustness

- Adding 5 % benign low-NIC aircraft (affected by something else) moves NLS from 26.0 to 42.5 km median (+64 %).
- GIST on the same scenes: 27.2 km median (p95 182.5): it learns to discount aircraft whose low NIC is unrelated to the emitter.
- Simulated jammers only; real-world localization is not yet validated.

INTERIM · SEED 0 · TEST

Ablations: What Each Component Buys

Variant	Question	Real F1	Real Spoof R@1%	Sim F1	Loc. median km
GIST (full)	—	0.88	0.83	1.00	18.4
w/o spatial graph attention	does sharing evidence across aircraft help?	0.86	0.74	0.99	16.5
integrity fields only	are physics features what catches spoofing?	0.75	0.29	0.82	18.5
w/o simulated training scenes	is the simulator needed for localization?	0.88	0.77	0.81	178.2



Graph: full → ablated

Real spoof R@1%: 0.83 → 0.74
Loc. p95: 164 km → 212 km
Neighbours catch spoofing and fix the worst localization misses.



Physics features: full → ablated

Real spoof R@1%: 0.83 → 0.29
Sim F1: 1.00 → 0.82
Spoofers fake integrity fields; physics features catch them.

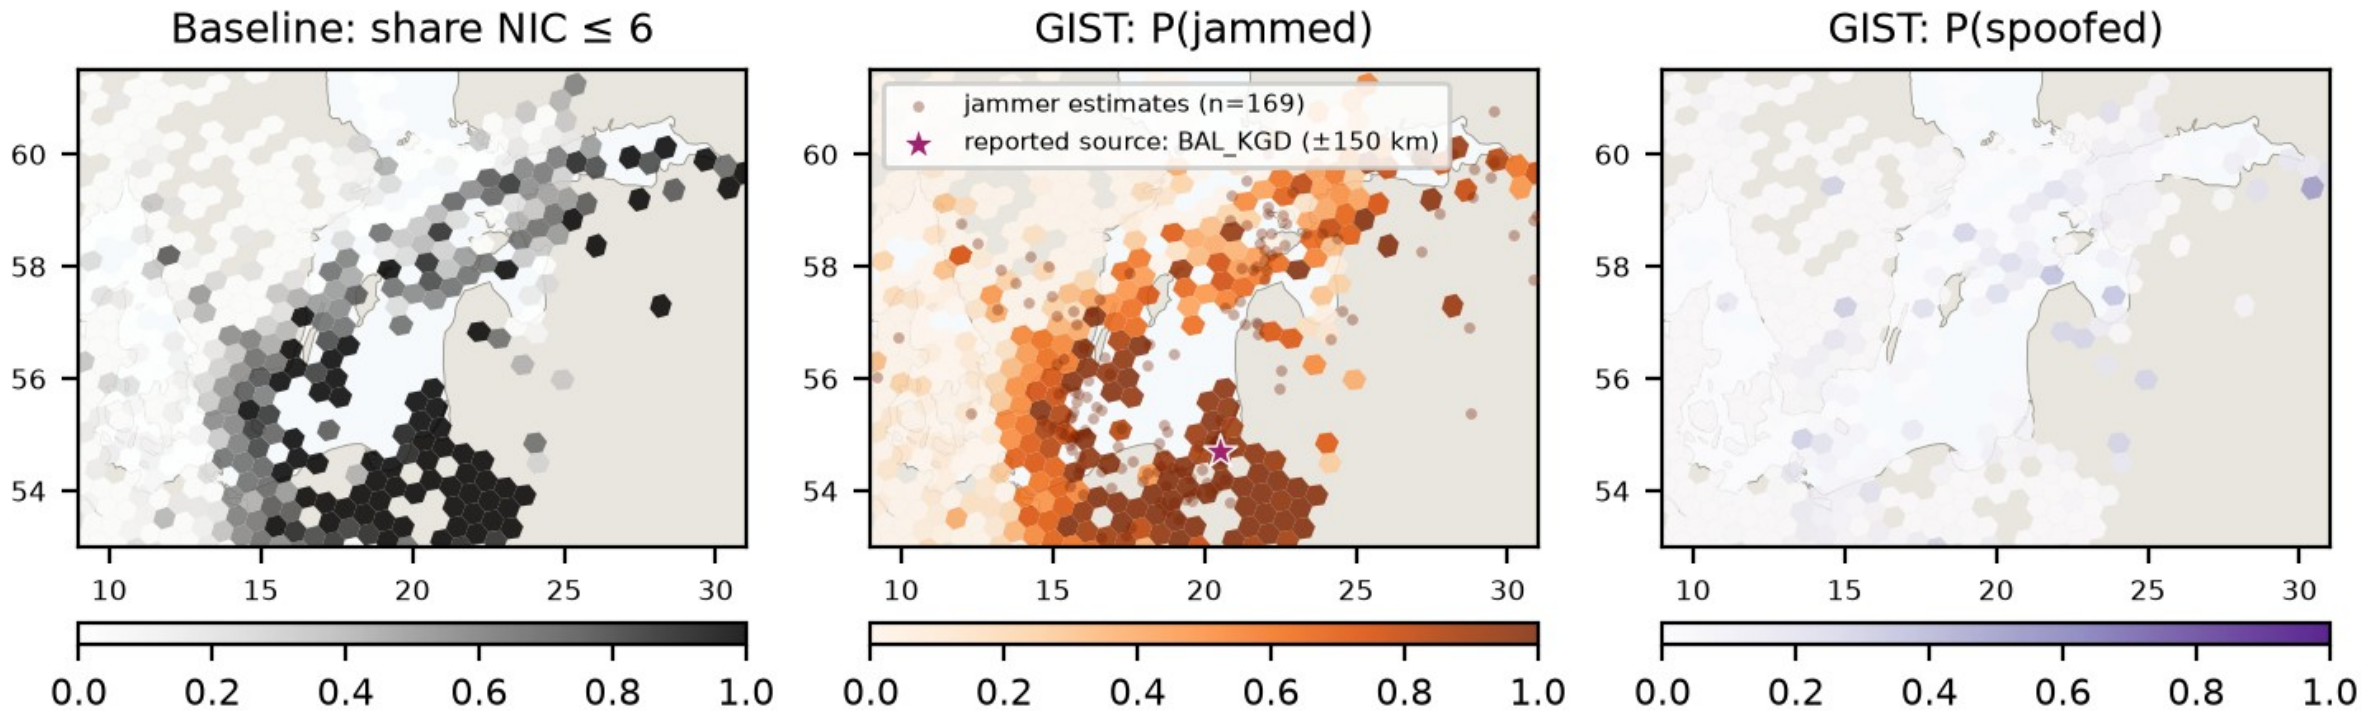


Simulation: full → ablated

Loc.: 18.4 km → 178.2 km
Sim F1: 1.00 → 0.81
Localization is learned almost entirely from simulation.

INTERIM MODEL · HELD-OUT TEST DAY

Case Study: Baltic Sea, 1 January 2026



Baseline
Share of aircraft with NIC ≤ 6 per hexagon: the gpsjam-style view.

GIST P(jammed)
The jamming footprint is centred on the reported Kaliningrad source (±150 km). The 169 jammer estimates are spread out: real-world localization is not validated yet.

GIST P(spoofed)
Near zero everywhere: consistent with the Baltic being a jamming region, not spoofing.

Challenges and Lessons Learned

Data licence

ADS-B Exchange samples forbid derivative works → switched to adsb.lol (ODbL).

Label-model artefact

All-abstain windows became “spoofed” (12.8 % in AUS) → left unlabelled.

Noisy wind

Naive wind triangle useless → fresh readsb wind, consensus per cell; feature, not LF.

8 GB RAM

MPS cache grew to 7.6 GB (epochs 2 → 36 min) → allocator watermarks; 21 × 21 grid.

Download throttling

GitHub assets fell to bytes/s after ~7 GB → resumable per-part fetch with back-off.

Flattering baselines

Clean-scene sampling helped NLS → added a benign low-NIC distractor split.

Lesson: check the data and the baseline before the model. Most of the time went into licences, labels and fair evaluation, and each fix changed the numbers more than any architecture change.

Limitations (Stated Honestly)

1	**Weak labels are not ground truth**	Real-traffic scores measure agreement with our physics assumptions. The MLAT LF is held out as an independent check; spoofed windows are few (374 in the 2026 test days).
2	**Simulation is simplified**	Free-space Friis loss, a coarse antenna model and a Gaussian receiver tolerance. Sim scores near 1.0 show the model learns the simulator, not that it is solved.
3	**Real-world localization unvalidated**	Localization numbers are on simulated jammers. The only real anchor so far (Kaliningrad) has ±150 km uncertainty; the Denver 2022 event with a known site needs OpenSky data.
4	**Uneven receiver coverage**	Crowdsourced receivers are sparse over parts of the Middle East and South Asia, so detection there is biased toward covered airspace.
5	**Interim numbers**	Results come from one seed trained on two days of data. The 34-day, 3-seed run is pending.

Conclusion

GIST shows that crowdsourced ADS-B carries enough physical redundancy to separate GNSS spoofing from jamming without a single hand-drawn label. Physics-consistency features expose spoofed aircraft that report high integrity; a Transformer with spatial graph attention turns them into per-aircraft classes; and a differentiable Friis / radio-horizon layer localizes simulated jammers more accurately than least squares (18.4 vs 26.0 km median on the same test scenes) and degrades less when unrelated low-integrity aircraft are present (27.2 vs 42.5 km).

Sees spoofing

Spoof recall @1% FPR 0.83 vs 0.01 for the NIC threshold.

No manual labels

10 physics LFs + Dawid–Skene; simulator for exact targets.

Physics-guided localization

Median 18.4 km on simulated jammers.

Numbers: interim seed-0 model on the 2026 test days; final 3-seed results pending.

Future Work

01

Real localization ground truth

Denver KDEN, Jan 2022 (FAA-documented jammer) via OpenSky Trino; access requested 2026-10-27.

02

More real events

Event catalogue for BAL, EMED, INPK with reported source areas and uncertainty.

03

Independent spoof checks

MLAT-only release and receiver radio-horizon plausibility as held-out evaluation signals.

04

Self-supervised pretraining

Masked-step reconstruction on unlabelled windows (head already in the model).

05

Live monitor

Minute-level latency evaluation and a static hex-map demo for a chosen date.

Source: PLAN.md sections 5.3, 6.3 and 8; DATA_LICENSE.md.

References

[1] Z. Liu, S. Lo and T. Walter, “GNSS interference source localization using ADS-B data,” Proc. ION ITM, 2022.

[2] Z. Liu, S. Lo, J. Blanch, Y.-H. Chen and T. Walter, “Locating GNSS interference sources using ADS-B with non-linear least squares,” NAVIGATION, vol. 72, no. 3, 2025.

[3] M. Schäfer et al., “Jammer on the Horizon: A robust method for GPS jammer localization using ADS-B data,” disco.cs.rptu.de.

[4] J. Wiseman, “GPSJAM GPS/GNSS interference map,” gpsjam.org.

[5] F. Eichstaedt, J. Budroweit and F. Stehle, “Detecting and measuring global GNSS signal interference via space-based ADS-B surveillance,” IEEE SHaRC, 2026.

[6] Aireon, “GNSS spoofing detection in aviation,” white paper v1.0, 2026.

[7] A. Ratner et al., “Snorkel: Rapid training data creation with weak supervision,” Proc. VLDB Endowment, vol. 11, no. 3, 2017.

[8] A. P. Dawid and A. M. Skene, “Maximum likelihood estimation of observer error-rates using the EM algorithm,” Applied Statistics, vol. 28, no. 1, pp. 20–28, 1979.

[9] S. Brody, U. Alon and E. Yahav, “How attentive are graph attention networks?,” ICLR, 2022.

[10] A. Vaswani et al., “Attention is all you need,” NeurIPS, 2017.

[11] ADSB.lol contributors, “adsb.lol globe_history open data (ODbL),” github.com/adsblol; Wiedehopf, “readsb,” github.com/wiedehopf/readsb.



Thank You

Questions and discussion

Data: ADSB.lol and contributors (ODbL-1.0) · Code: MIT

Vansh Agarwal (1BI23IS133) · Vichitra Attri (1BI23IS135) · Vidya Bai K (1BI23IS137)

Faculty: Dr. Asha T., Deep Learning course, Department of ISE, Bangalore Institute of Technology