Hide and Stream: Ambient Aircraft-Doppler Geolocation of Live Audio, and a Defense

Remy Ochei

Draft — November 2026

Abstract

Aircraft passing overhead leave slowly drifting tonal lines in the background of a recording — fan and blade-pass tones, often below the threshold of conscious hearing. The *pitch drift* of such a line is the acoustic Doppler signature of a single flyover, and it fits a four-parameter physical model (emitted frequency, ground speed, slant miss distance, time of closest approach) once propagation delay is accounted for. Because public ADS-B networks publish the real, worldwide tracks of the aircraft that produced those lines, an adversary with only a published recording and open flight data can score candidate listener positions by how well each explains the observed drifts, localizing the recording toward a city block. Conventional noise reduction and high-pass filtering do not remove the signature, because it is a drift, not a band. We characterize the leak, give an open, local-only detector (`contrail`) that quantifies a clip's exposure, and demonstrate accuracy against known-anchor clips containing no person — corroborated by a single error-only self-test on the author's own pooled recordings, whose location is never published. We then present a scrubber that destroys the signature at low perceptual cost (voice-band correlation > 0.9) and a decoy renderer (`tsukuyomi`) that borrows *real* overflight tracks as defensible chaff. We release the defenses and the physics, disclose the attack only at the level needed to motivate them, withhold any locator, and notify platform creator-safety teams ahead of publication.

1. Introduction

Creators who take every known precaution against *visual* location leaks — camera off, windows shut, metadata stripped, no landmarks on screen — still broadcast their rooms. And a room is not acoustically silent about where it is. When an aircraft passes overhead, its tonal noise reaches the microphone as a narrowband line whose frequency drifts downward through closest approach in a shape fixed by physics. That drift encodes the geometry of the pass relative to the microphone; matched against the public record of which aircraft were actually in the sky, it encodes the microphone's place on Earth.

This paper makes that leak concrete and then defends against it. Our stance is deliberately asymmetric: we believe the capability already exists in private hands, that the affected population is large and undefended, and that the responsible response is to ship the defense while disclosing only enough of the attack to justify it.

Contributions. 1. **Characterization.** We describe the aircraft-Doppler leak in speech-dominated audio and show why it survives the filters creators already use (§3). 2. **An open detector.** `contrail` finds the drifting lines, fits each to the retarded-time flyover model, gates on physical plausibility, and emits a per-clip *leak score* — running entirely on the creator's machine, audio never leaving it (§4). 3. **A victim-free evaluation.** Accuracy is measured against synthetic and chosen-anchor clips that contain no person, and corroborated by a single error-only real-world self-test on the author's own pooled recordings (§5). 4. **A defense.** A scrubber that destroys the signature at low voice cost, and `tsukuyomi`, which renders *real* overflight tracks as chaff a flight-data cross-reference cannot

discard for free (§6). We include a measured negative result on naive VAD gating. 5. **Responsible disclosure.** We publish the defenses and the physics, withhold the locator, and coordinate with platform creator-safety teams and a creator OPSEC guide (§7).

Scope. This paper and its released code contain a detector, a scrubber, a decoy renderer, and the physics. They do not contain, and we do not publish, any matcher or listener-grid scorer that turns a third party's recording into coordinates.

2. Threat model

Adversary. Passive and remote. The adversary has only (a) media the target published — a VOD, clip, or upload containing background audio — and (b) public ADS-B flight history from volunteer networks (e.g. adsb.lol, ODbL, no account). No privileged access: no malware on the target's machine, no network-path observation, no platform insider data, no interaction with the target. This is the weakest-capability adversary that still succeeds, which is what makes the leak serious: the raw materials are a public broadcast and an open data set.

Target. Any creator who broadcasts or uploads audio recorded in a fixed location with overflight exposure and without delay or masking: live streamers, IRL/VTuber broadcasters, podcasters, anyone whose ambient room audio reaches an audience.

Goal. Recover the recording location. Detection (§4.1–4.2) recovers each flyover's parameters relative to the microphone; localization (§4.3) converts a set of passes plus public flight history into a position estimate, resolving toward a city block as the number of usable passes grows.

Assumptions and preconditions. - Aircraft actually overflew during the recording and broadcast ADS-B/MLAT (not all do). - Receiver coverage exists for the area and epoch; volunteer coverage is uneven and history is patchy, so an empty result does not prove a clean sky. - The recording is long enough, and SNR high enough, to recover several passes; a single pass constrains geometry weakly. - Timing is resolvable to ~10 s (the ADS-B heatmap slice).

Out of scope. Active or network-level deanonymization, visual leaks (treated as related work, §8), and targets in airspace with no ADS-B coverage. Mobile recordings and curved/climbing tracks degrade the straight-pass fit (§9).

Defender. The same creator, acting locally and before publishing: audit a clip's exposure and scrub it (§6), or adopt behavioral/platform mitigations (§6.4). The defender never needs to send audio anywhere.

3. The physics of the leak

3.1 What radiates

An aircraft in cruise or approach emits strong narrowband tonal lines — fan/blade-pass tones and their harmonics, typically tens to a few hundred Hz. To a fixed microphone on the ground, a single straight, level, constant-speed pass produces a characteristic *falling-S* frequency track: nearly flat while the aircraft approaches from far off, a steep monotonic drop through closest approach, flattening again as it recedes. The received frequency is the classical Doppler relation

$$f_{obs}(t) = f0 \cdot c \,/\, (c - v_{rad}(t)), \quad c = 343\ \mathrm{m/s},$$

where v_{rad} is the radial velocity of the source toward the listener and $f0$ the emitted tone.

3.2 Why the curve is informative

For a straight level pass, the geometry relative to a listener at the origin is fully described by four parameters: the emitted frequency **f0**, the ground speed **v**, the slant miss distance **d** at closest approach, and the time of closest approach **tc**. Fitting an observed track to this model recovers all four. Critically, these are recovered *relative to the microphone* — they describe the pass, not the microphone's position on Earth. Recovering position is a separate step (§4.3).

3.3 Retarded time is not optional

Sound is slow. The frequency heard at observation time t was emitted at an earlier *emission time* te, where $t = te + r(te)/c$ and r is the source–listener range at emission. Solving for te is a quadratic:

$$(c^2 - v^2)\cdot te^2 + (-2c^2 t + 2v^2\cdot tc)\cdot te + (c^2 t^2 - v^2\cdot tc^2 - d^2) = 0,$$

taking the root with $te < t$. Observed frequency then uses the radial velocity $v_rad = -v\cdot x/r$ with $x = v(te - tc)$, $r = \sqrt{(x^2 + d^2)}$. Ignoring propagation delay visibly warps the curve: at a slant distance of ~2.5 km the one-way delay is ~7 s and varies by several seconds across a single pass, shifting the apparent time of closest approach and biasing every parameter. The model (`contrail.doppler.Flyover`, `contrail.detect.doppler_model`) carries the delay exactly.

3.4 Why ordinary filtering misses it

The signature is a *pitch drift*, not energy in a fixed band. High-pass filtering, spectral subtraction, and generic denoisers all operate on frequency bands or stationary noise estimates; a line that slides a few Hz through the low band over a minute passes straight through them. This is the core reason the leak survives normal "clean-up."

4. Detection (`contrail`)

`contrail` is the open detector. It answers one question — *does this clip contain recoverable flyover lines, and how exposed is it?* — entirely on the user's machine, with no network and no location data.

4.1 Front end

- **Long-window STFT.** A 1 s Hann window (→ ~1 Hz bins), hop 0.25 s, over the 50–800 Hz band. Long windows are deliberate: aircraft lines drift at a few Hz/s and stay sharp, while speech harmonics jump tens of Hz between syllables and smear out. The window length is what separates the two.
- **Per-frame spectral whitening.** A running median along frequency (25 bins) is subtracted per frame, so a line 20+ dB below the voice bands survives instead of being masked.
- **Sub-bin peak picking.** Peaks above a 6 dB prominence, refined to sub-bin frequency by parabolic interpolation.

4.2 Tracking → physical gate

1. **Greedy ridge tracking,** downward-only: a receding aircraft never rises in pitch, so links are accepted only within a slope-predicted tolerance (~2.5 Hz, widened across short gaps) and never upward. Short dropouts (≤4 frames) are bridged for speech-harmonic crossings.
2. **Fragment merge, then split.** Speech harmonics sweeping the 200–400 Hz band steal peaks and fragment a line right at its steepest point; fragments are rejoined when a *joint* fit stays tight. Conversely, two crossing aircraft lines can be glued by the greedy linker; an over-long track

whose fit is poor is split at the interior point minimising the two half-fits, and the halves re-queued.

3. **Physical fit as the gate.** Every surviving candidate is fitted to the four-parameter retarded-time model by Levenberg–Marquardt (`fit_doppler`, analytic central-difference Jacobian). A candidate is kept only if the residual is < 1 Hz RMS *and* the fitted speed (40–320 m/s) and slant distance (0.1–12 km) are physically plausible. This physical gate — not a smoothness heuristic alone — is what suppresses false positives from quasi-tonal noise.

A per-clip **leak score** aggregates the surviving tracks into a single exposure number a creator can act on. On the grouping bench (8 seeds), line-to-plane pairing achieves **precision 16/16, recall 16/23, with zero off-truth lines** — i.e. the physical-fit gate admitted no spurious lines, at the cost of some under-merging (missed joins), consistent with recall below precision.

4.3 The attack, at disclosure level

Detection recovers each pass's parameters relative to the microphone. Localization is a *separate* step and is where the privacy loss occurs: given the detected tracks and **public** ADS-B flight history (volunteer networks such as adsb.lol publish worldwide tracks at ~10 s resolution under an open licence), an adversary scores candidate listener positions by how well the real aircraft that actually flew during the recording — their true tracks, times, and speeds — would reproduce the observed Doppler lines at that position, against control scores to reject coincidence. The position whose real overflights best explain the recording localizes it; with enough passes this resolves to roughly a city block.

We describe the attack at this level deliberately. It is sufficient to establish that the threat is real, reproducible from public data, and not mitigated by conventional audio cleanup. We do not release the matcher, a listener-grid scorer, or any artifact that turns a third party's recording into coordinates. The demonstrations in §5 use flight data over anchors we choose, with no real person as subject (`tsukuyomi`, §6.2).

5. Demonstration (no human subject required)

Because localization accuracy can be measured against *known* ground truth, the evidence needs no victim:

- **Substrate.** `tsukuyomi` renders the aircraft that genuinely flew over an anchor *we* choose (real hex codes, real tracks, real times, from public ADS-B) into Doppler audio, with a synthetic room bed and **no voice/person**. The anchor is a place, tied to no one.
- **Detection metrics.** Recall/precision of `contrail` on these clips and on synthetic-planted benches, plus parameter error (speed, miss distance, f0) vs. truth.
- **Localization error.** Run against clips with a *known* anchor and report distance error — this quantifies the attack's accuracy without locating any person.
- **Real-speech robustness.** Mixed with real speech to measure degradation; see the VAD negative result in §6.3.

5.5 Real-world validation (author self-test, error-only)

To confirm the method transfers from chosen-anchor clips to genuine in-the-wild audio, the author ran a self-test over several of his own recordings, at a location he is entitled to disclose (his publicly listed business address).

- **What is reported:** the localization *error* only — the recovered position's distance from the ground truth the author holds. On the pooled self-test audio the method localized to within **227 m** of truth — a single grid cell. The estimate was built from 23 usable Doppler tones across roughly 135 hours of the author's own pooled recordings, 6 of them clock-consistent at the recovered cell.
- **What is never reported:** the recovered position, the ground truth, the recording date/place, or anything that narrows them — at any granularity, exact or city-block. The coordinates do not appear in the paper, the artifact, or the figures.
- **Why error-only.** §5.1–5.4 already establish accuracy against known anchors with no human subject; the real-world test adds only the fact of transfer to genuine field audio. Publishing the author's actual location would add exposure, not evidence, so we withhold it by design.
- **Third parties.** The recordings were made where no non-consenting person is implicated by the result; because only the error is published, no co-resident or location trail is exposed regardless.
- **Consent.** The only subject of this test is the author, on his own recordings and at a location he is entitled to disclose; even so, only the error is reported, never the location. No other person was located, with or without consent, for any result in this paper (see §7).

6. The defense

6.1 Scrubbing (`contrail.scrub`)

Three STFT-domain tiers, verified by re-detection (what the scrubber edits is exactly what a detector sees):

- **gate** — in no-voice frames, attenuate < 500 Hz by 40 dB. Weak alone when pauses are short.
- **notch** — for each detected ridge and its harmonics, zero a ±2.5 Hz band that follows the ridge frame-by-frame. This does the real work (synthetic bench: 1/9 planted lines survive), but only on lines the detector *found* — scrub efficacy is capped by detector recall, the argument for shipping detector and scrubber together.
- **decoy** — ambiguity rather than removal (see 6.2).

Current bench (regenerated, 9 planted lines across 3 clips):

tier	real lines still detectable	decoys detected
gate only	9/9	—
notch only	**1/9**	—
decoy only	7/9	13/24
all three	**1/9**	17/24

Voice cost is low: notched audio correlates > 0.9 with the input in the voice band; notches are a few Hz wide, below the formants.

6.2 Honest chaff (`tsukuyomi`)

Random synthetic decoys are discarded for free by an adversary cross-referencing flight history, because they match no real aircraft. `tsukuyomi` instead mixes in the Doppler of planes that *actually* flew over some unrelated anchor — real hex, real tracks — so a flight-data cross-reference cannot trivially separate chaff from the true line. This raises the matcher's ambiguity with real, defensible decoys.

6.3 A negative result worth reporting

A naive VAD band-gate — suppressing low-frequency peaks in talking frames — was validated on ~48 min of real speech (3 recordings, 12 windows) with planted flyovers and found **not** to help at any margin: a low margin fires constantly and deletes the aircraft lines that share the speech band; a high margin barely fires. When speech is near-continuous you cannot blanket-suppress the band the line occupies. A real fix needs a learned VAD (silero/webrtcvad) with *targeted* rejection of peaks coincident with a tracked voice harmonic, not a band gate. (Reported honestly because it bounds what the cheap defense can do.)

6.4 Mitigations beyond the tool

Broadcast delay, a masking audio bed, closed windows, directional mics; and at platform scale, delay defaults and bundled scrubbing — the highest-leverage fix, since the affected population is aggregated at the platform.

7. Ethics and responsible disclosure

This work studies a privacy attack in order to defend against it. We take the ethics of that posture seriously and state our choices explicitly, because the line between "privacy disclosure" and "attack release" is exactly where papers like this live or die.

7.1 Human subjects and consent

Every localization *result* in this paper comes from one of two sources, both consensual: (a) synthetic and chosen-anchor clips that contain no person (§5.1–5.4, via `tsukuyomi`), and (b) a single real-world self-test on the author's own pooled recordings, at a location the author is entitled to disclose, reported error-only (§5.5). No person was located without consent for any published result. We report the real-world test as an error metric only; no subject's recovered location, exact or coarse, appears anywhere in the paper, artifact, or figures.

7.2 A vulnerability encountered in the wild

The authors became aware of this leak through an instance in which a creator was localized from public material. That instance is **not** the evidence here: its data, identity, figures, and location are excluded entirely and are not reproduced, and the affected party was not re-contacted for this work in any way that would further expose them. We name it only to be candid about provenance and to be unambiguous that non-consensual localization is precisely the conduct this paper exists to prevent. A finding's origin does not license its publication; consent governs what we print.

7.3 What we release, and what we withhold

- **Released (defense):** `contrail` (detector + leak score + scrubber) and `tsukuyomi` (honest chaff + no-person demo clips). These let any creator audit and scrub their own audio locally; audio never leaves the machine.
- **Released (attack), disclosure-level only:** the physics (§3) and the feasibility of ADS-B correlation (§4.3) — enough to establish that the threat is real, reproducible from public data, and unmitigated by conventional cleanup.
- **Withheld:** the locator engine — any matcher, listener-grid scorer, or pipeline that turns a third party's recording into coordinates. No such artifact is published, and the defensive repos are deliberately scoped to exclude it. We ask that forks keep it that way.

The asymmetry is intentional: we publish what lets people defend, and the minimum about the attack needed to motivate that defense — not a build guide.

7.4 Coordinated disclosure

Because the leak is environmental rather than a defect in any one product, there is no single vendor to patch it; the highest-leverage mitigation sits with the platforms that aggregate affected creators. Ahead of publication we notify major streaming platforms' creator-safety / Trust & Safety teams with the class-level advisory and the mitigations (delay defaults, masking, bundled scrubbing), so protection can reach the population before the technique is public. We also publish a creator-facing OPSEC guide (companion) describing the risk class and defenses with no target and no method.

7.5 Residual risk

An adversary could reconstruct a locator from public ADS-B data and the physics described here; we cannot prevent that and do not claim to. Our argument is that defenders are currently *unaware* and *undefended*, that the detector and scrubber materially change that, and that shipping the defense plus disclosure-level attack detail does more good than withholding everything while the capability already exists in private hands.

8. Related work

Electric Network Frequency (ENF). The closest prior channel, and the sharpest contrast. Mains-powered recordings pick up the grid's ~50/60 Hz hum and its small, grid-wide frequency fluctuations. ENF forensics uses this chiefly for *timestamping* (aligning a recording's hum trace to a logged grid-frequency reference), reliably for *grid/interconnect identification* (continental-scale), and — contested and noisy — for coarse *intra-grid* location. The decisive difference from our channel is not resolution but **reference-data availability.** ENF localization or timestamping requires a continuous, location- or grid-tagged *reference trace* of grid frequency; there is no open, global, historical grid-frequency feed — such references are largely held by forensic labs, utilities, and research groups, and coverage is national and partial at best. Our channel's reference layer is the opposite: ADS-B flight history is public, free, worldwide, and open-licensed (adsb.lol, ODbL), with no account. ENF is therefore mostly a *forensic-lab* capability against a recording in hand; aircraft-Doppler is an *open-adversary* capability usable by anyone with the published clip. Secondarily, the carrier differs in kind — a *moving external source* whose Doppler geometry, matched to specific aircraft kinematics, yields a *metric* city-block estimate rather than a grid region, and requires no mains power, so it reaches the battery/IRL recordings ENF cannot touch.

Visual geolocation. A large body of work localizes images/video from content: EXIF geotags, skyline and landmark matching, sun position and shadow geometry, reflections, and learned image-to-GPS retrieval (e.g. IM2GPS-style methods and their successors). These are the dominant media-geolocation threat and the reason visual OPSEC is well known. The audio channel is complementary and under-appreciated: it leaks through a closed-window, camera-off, voice-only stream where every visual mitigation already holds.

Acoustic flyover parameter estimation (the detector's prior art — stated plainly). Recovering an aircraft's speed, altitude, and source (blade-rate) frequency from the Doppler drift of a single ground microphone is a mature technique in aircraft-noise monitoring: least-squares fits of the blade-rate tone's time variation, short-time Fourier / Wigner–Ville instantaneous-frequency tracking, polynomial-phase (DPT) methods that add the amplitude cue, and the single-sensor Lloyd's-mirror ground-reflection approach. **We claim no novelty in the detection physics.** Our contributions over

this body of work are (a) making it run on *covert, speech-dominated consumer stream audio* rather than clean monitoring recordings of propeller aircraft — hence the whitening, downward-only tracking, and physical-fit gate that survive a talking human in-band — and (b) the privacy framing: the recovered per-pass parameters are not the goal but the input to localization.

Passive acoustic localization / multilateration. Classical acoustic source localization and TDOA multilateration use *arrays* of known receivers to locate an unknown source; Doppler-mesh methods localize aircraft from several distributed mics given partial prior kinematics. We invert that geometry: a *single, unknown* receiver is localized by matching known moving sources (aircraft) against their *public* tracks. The inversion — unknown listener, known sources, public reference — is the privacy-relevant novelty, not the acoustics.

ADS-B as a data source. A line of work studies ADS-B integrity, spoofing, crowdsourced receiver networks (OpenSky), and coverage. We use ADS-B only as a public, benign reference; the novelty is the acoustic side of the join, not the flight data. The same openness that enables these community networks is what makes the reference layer of our attack freely available (contrast ENF, above).

References

ENF - Ngharamike et al., "ENF Based Digital Multimedia Forensics: Survey, Application, Challenges and Future Work," *IEEE Access* 11:101241–101272, 2023. DOI 10.1109/ACCESS.2023.3312181. - Hua, Bi, Thing, "On Practical Issues of Electric Network Frequency Based Audio Forensics," *IEEE Access* **5:20640–20651**, 2017. DOI 10.1109/ACCESS.2017.2647921. - Maiberger, Gusakov, Routtenberg, "Empirical Evaluation of ENF Extraction Methods for Accurate Timestamping in Multimedia Forensics," *CISS 2025*. DOI 10.1109/CISS64860.2025.10944719. (Ben-Gurion Univ.)

Image geolocation - Hays & Efros, "IM2GPS: Estimating Geographic Information from a Single Image," *CVPR* 2008. DOI 10.1109/CVPR.2008.4587784. (Benchmark: ~16% of 237 photos within 200 km.) - Weyand, Kostrikov, Philbin, "PlaNet — Photo Geolocation with Convolutional Neural Networks," *ECCV* 2016. arXiv:1602.05314. - Vo, Jacobs, Hays, "Revisiting IM2GPS in the Deep Learning Era," 2017. arXiv:1705.04838.

Acoustic flyover estimation *(the detector's prior art — §8)* - Ferguson & Quinn, "Application of the short-time Fourier transform and the Wigner–Ville distribution to the acoustic localization of aircraft," *J. Acoust. Soc. Am.* **96(2, Pt.1):821–827**, Aug 1994. Foundational single-mic model: turbo-prop blade-rate via STFT/WVD → speed, altitude, source frequency; 18 transits (150/200/250 kn × 250–1500 ft). - Reid, Zoubir & Boashash, "Aircraft flight parameter estimation based on passive acoustic techniques using the polynomial Wigner–Ville distribution," *J. Acoust. Soc. Am.* **102(1):207–223**, July 1997. (constant height, ground speed, range, acoustic frequency from a single stationary observer; PWVD as the IF estimator; optimal window-length analysis.) - Lo, Perry & Ferguson, "Aircraft Flight Parameter Estimation Using Acoustical Lloyd's Mirror Effect," *IEEE Trans. Aerospace and Electronic Systems* **38(1), Jan 2002, p. 137–**. Single-sensor; direct/ground-reflected interference fringes; NLS + generalized Hough estimators, CRLB analysis; field accuracy ±22.6 kn speed, ±126.2 ft altitude over 10 transits. - Martín Román, *Passive acoustic method for aircraft localization*, PhD thesis, Universitat Politècnica de Catalunya (UPC), 2013. Journal version: "Passive acoustic method for aircraft states estimation based on the Doppler effect," *IEEE Trans. Aerospace and Electronic Systems*, 2014. (Doppler among spatially distributed mics; ambiguity-function frequency-stretch between mic pairs; genetic-algorithm position/velocity estimate; ~7 mics + prior kinematics at one instant; works for jets, not just low props; RC-aircraft field test.)

ADS-B / OpenSky - Schäfer, Strohmeier, Lenders, Martinovic, Wilhelm, "Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research," *IPSN* 2014, Berlin, pp. 83–94. DOI 10.1109/IPSN.2014.6846743. - Strohmeier, Lenders, Martinovic, "Lightweight Location Verification in Air Traffic Surveillance Networks," *CPSS '15* (1st ACM Workshop on Cyber-Physical System Security), Singapore, pp. 49–60, 2015. (TDOA-based verification; Best Paper.) See also their survey "On the Security of the ADS-B Protocol," *IEEE Comms Surveys & Tutorials* 17(2):1066–1087, 2015. - Leonardi, Strohmeier, Lenders, "Jamming/Garbling assessment and possible mitigations in the OpenSky network," *Proc. 7th OpenSky Workshop*, pp. 85–96, 2019. - Darabseh, Khandker, Pöpper, "Securing the Sky: Detecting Aircraft Location Drifting through Cross-Checking Receiver-Based Estimated and Received ADS-B Trajectories," *Journal of Open Aviation Science*, 2023/2024. (Uses receiver positions, not precise TOA — suits crowdsourced OpenSky.)

Audio OPSEC / anti-doxxing. Practitioner guidance on streamer safety covers visual leaks and metadata thoroughly but rarely ambient-audio geolocation beyond "mute notifications." We contribute a concrete, measured instance of the audio risk and an actionable defense.

9. Limitations and future work

- **Real-world recall is only partly characterized.** Detector performance rests on synthetic and chosen-anchor benches plus the author's pooled self-test recordings (§5.5) and ~48 min of real speech for the VAD study (§6.3). Recall across diverse rooms, mics, codecs, and airspaces is future work.
- **Straight-pass model.** The fit assumes a straight, level, constant-speed track. Real tracks curve, climb, and change speed; `tsukuyomi` already renders the true curved tracks, and generalizing the *detector's* fit beyond the straight model is open (it currently tolerates this as fit residual, dropping the worst cases).
- **Airspace dependence.** Remote or low-coverage areas yield few usable passes; the attack weakens exactly where few planes fly, and the defense's decoys depend on there being real traffic to borrow.
- **Scrub capped by recall.** Notching only removes lines the detector finds (§6.1); an undetected line survives. Improving recall is the single highest-value detector work.
- **Adaptive adversary.** A matcher aware of `tsukuyomi` chaff may try to down-weight borrowed tracks; the chaff/counter-chaff dynamic is unquantified.
- **Codec and delay interaction.** How lossy streaming codecs and platform transcoding attenuate or preserve the low-frequency lines, and how a few minutes of broadcast delay interacts with history-based matching, both merit measurement.

10. Conclusion

A room is not silent about where it is. Aircraft overhead write a physical, public-referenced signature into ambient audio that survives the filters creators already trust, and matching it against open flight data localizes a recording to roughly a city block — through a camera-off, window-closed, voice-only stream where every visual precaution already holds. We characterized the leak, shipped an open local-only detector and scrubber, demonstrated accuracy without locating any person, and disclosed the attack only as far as needed to motivate the defense while withholding the locator. The sky is an audible reference signal; the responsible response is to make the defense as freely available as the attack's raw materials already are.

Acknowledgements

The author thanks Sylvia Vale and Mel De Leon (meldeleon.com) for discussion, encouragement, and feedback during this work. Any errors are the author's own.