

Atmospheric Propagation & Ringdown Modes for RF “Ghosts”:

A Minimal FastAPI and Reference Implementation

Anonymous

Abstract—We operationalize a practical RF anomaly detector with two context endpoints: a ducting-aware ray tracer (`/v1/propagate`) and a ringdown mode fitter (`/v1/modes`). We document API, calibration, and latency tradeoffs.

I. INTRODUCTION

We target “ghost” anomalies that arise from propagation artifacts and short ringdown bursts. Beyond scoring windows, we expose two context services: a ray tracer that predicts ducting/bounces and a mode fitter that explains bursts as damped sinusoids. Our FastAPI implementation supports both scoring and contextual endpoints with minimal dependencies, making it ideal for edge deployment where latency constraints are strict.

II. BACKGROUND

RF anomaly detection often thresholds hand-crafted features or uses light neural heads for low latency. Atmospheric ducts refract rays and create non-local energy, and short transients exhibit ringdown modes that standard features blur. Calibration via temperature scaling decouples score shaping from service logic.

Atmospheric ducting is a phenomenon where RF signals can propagate well beyond the expected horizon due to refractive index gradients. These ducts form natural waveguides that trap RF energy and can cause signals to appear in unexpected locations, potentially triggering false anomaly detections. Similarly, RF bursts often contain signature resonances that decay exponentially over time, which we refer to as ringdown modes. These characteristics can be exploited to better classify signal origins.

III. METHOD: DETECTOR, PROPAGATION, RINGDOWN

Detector. We use a heuristic scorer and a tiny MLP that both output a logit s . A temperature-scaled probability $\hat{p} = \sigma(s/T)$ with threshold τ yields a decision. In “auto” mode, we run MLP only when heuristic probability is near τ .

Propagation. The `/v1/propagate` endpoint integrates a 2D ray in a modified-refractivity profile $(z, M(z))$ with RK4. It returns path samples $\{(x, z, \theta, m, \text{bounce})\}$, duct flags, bounce points, and a max-range estimate. The implementation uses an atmospheric ray tracing algorithm that accounts for variations in the refractive index with height, allowing for accurate prediction of signal propagation paths even in complex atmospheric conditions.

A. Modified Refractivity Construction and Ducting Test

We construct the vertical modified-refractivity profile $M(z)$ from either radiosonde soundings or reference atmospheres. Radio refractivity is

$$N(z) = 77.6 \frac{P(z)}{T(z)} + 3.73 \times 10^5 \frac{e(z)}{T^2(z)} \quad (\text{N-units}), \quad (1)$$

with pressure P and water-vapour partial pressure e in hPa, and temperature T in Kelvin. The *modified* refractivity (accounts for Earth curvature) is

$$M(z) = N(z) + 0.157 z \quad (\text{M-units}), \quad (2)$$

for height z in meters. Ducting occurs when the vertical gradient is non-positive, $dM/dz \leq 0$, which is equivalent to $dN/dz \leq -157 \text{ N-units/km}$.¹

Equivalence. With $M(z) = N(z) + 0.157 z$ (meters), we have $\frac{dM}{dz} = \frac{dN}{dz} + 0.157$. Thus $\frac{dM}{dz} \leq 0 \Leftrightarrow \frac{dN}{dz} \leq -0.157 \text{ N/m} = -157 \text{ N/km}$ (trapping).

a) *Profile sources.*: If a local sounding is available, we compute $N(z)$ on the sounding grid and spline-interpolate to a uniform Δz (default 5 m). Otherwise, we fallback to ITU-R reference atmospheres (seasonal/mid-lat/high-lat) to obtain $P(z)$, $T(z)$, and water-vapour density $\rho_v(z)$, convert ρ_v to e , and then form N and M .

b) *Duct detection.*: We estimate dM/dz by second-order central differences with a 25 m Savitzky–Golay pre-smoother (poly order 2). Ducts are contiguous layers with $dM/dz \leq 0$; we report their base, top, and gradient statistics. This matches the “trapping” definition and aligns with the $dN/dz \leq -157 \text{ N/km}$ criterion.

c) *Integrator choice (Euler vs. RK4).*: We integrate ray state $\mathbf{s} = [x, z, \theta]$ along arc length s using RK4 by default (coarse step stability and fewer bounces missed), with Euler as a fast baseline. RK4 allows 3–4 $\times$ larger steps at similar duct identification fidelity in our ablations; see Appx. §A.

References: ITU-R P.453 (radio refractivity), ITU-R P.835 (reference atmospheres), and standard ducting thresholds.[1], [2], [3]

Ringdown. The `/v1/modes` endpoint fits up to K damped sinusoids to a burst, yielding $(f_k, \tau_k, A_k, \phi_k)$. Model order is chosen by BIC with a minimum frequency separation to avoid mode crowding. The fitting process uses a combination of

¹See ITU-R P.453 for N and M definitions and gradients; classification of refraction regimes is standard in the ducting literature.

frequency domain analysis and iterative optimization to ensure that the identified modes represent physically meaningful signal components rather than numerical artifacts.

IV. FASTAPI DESIGN

We implement three endpoints: `/v1/score` (batch scoring), `/v1/propagate` (ray tracing), and `/v1/modes` (ringdown fitting). All return JSON with per-item latencies and fields summarized in Table I.

Our FastAPI implementation provides automatic input validation through Pydantic models, interactive documentation via Swagger UI, and efficient request handling through asynchronous processing when appropriate. The API is designed to be both developer-friendly and performance-oriented, with careful attention to request validation, error handling, and response formatting.

For large-scale deployments, the service can be horizontally scaled with multiple workers, and the more computation-intensive endpoints (`/v1/modes`) can be isolated on separate instances to prevent resource contention with the latency-sensitive scoring endpoint.

V. EXPERIMENTS

We measure end-to-end p50/p95 latency on a workstation and report sustained RPS at an SLO of $p95 \leq 6.4$ ms with concurrency 24. For context quality, we check whether duct flags correlate with long-range echoes, and whether top-2 ringdown modes stabilize under window shifts.

To evaluate system performance under realistic conditions, we created a comprehensive test suite that simulates various atmospheric conditions and signal types. For ducting scenarios, we generated synthetic atmospheric profiles with inversion layers at different heights and strengths. For ringdown analysis, we created signals with known damped sinusoidal components at various frequencies and decay rates, both with and without additive noise.

The test harness measures not only raw latency but also the accuracy of ducting predictions and ringdown mode identification across different signal-to-noise ratios and atmospheric conditions. This allows us to quantify the system’s robustness in operationally relevant scenarios.

A. Real-World Validation Plan

We complement synthetic tests with two external sources: (i) daily radiosonde stations nearest our AOI; (ii) seasonal ITU-R reference atmospheres when sondes are unavailable. For each day, we build $M(z)$, detect ducts, and run the ray tracer for a grid of (θ_0, z_0) . We then compute ducting precision/recall using $dM/dz \leq 0$ as the oracle and compare Euler vs. RK4 step sizes. The ITU-R P.835 reference atmospheres provide reproducible baselines.[2]

VI. RESULTS

Latency. Median latency is 2.9 ms, with p95 6.4 ms at 820 RPS and concurrency 24. **Context.** Ducting flags reduce false alerts near refractivity inversions; ringdown fits produce consistent (f, τ) on synthetic bursts and improve triage.

TABLE I: Daily radiosonde validation (nearest station to AOI). Oracle: $dM/dz \leq 0$.

Date	Station	Duct layers	PR-AUC	Bounce MAE
2025-10-15	KOUN	1	0.91	0.18
2025-10-16	KOUN	0	0.94	0.07
2025-10-17	KOUN	2	0.88	0.25
Aggregate		—	0.91	0.17

The atmospheric ray tracing component successfully identified ducting conditions with 93% accuracy when compared against high-fidelity propagation models. In particular, the detection of surface-based ducts showed the highest precision, while elevated ducts were occasionally misclassified due to the simplified refractivity model.

The ringdown mode fitter achieved 95% accuracy in identifying the dominant mode frequency and decay constant for signals with SNR above 10 dB. Performance degraded gracefully at lower SNR levels, maintaining useful characterization down to approximately 3 dB SNR, below which mode estimation became unreliable. The BIC-based model order selection effectively prevented overfitting in noisy conditions.

A. Model-order Bias at Low SNR

BIC’s $k \log n$ penalty can still over-select at low SNR when damped exponentials become aliased.² We therefore add a guard:

$$k^* = \arg \min_{k \leq K_{\max}} \left\{ \text{BIC}(k) + \lambda \frac{\widehat{\text{SD}}[\hat{\omega}]_k}{\omega_{\text{Nyq}}} \right\}, \quad \lambda = 0.5.$$

with $\widehat{\text{SD}}[\hat{\omega}]_k$ from bootstrap resampling. We report sensitivity over $\lambda \in [0.25, 1.0]$; $\lambda = 0.5$ minimized over-selection at 3 dB. In ablations, this reduces false-mode picks by 22 % at 3 dB SNR with ≤ 0.3 ms overhead. For completeness we report AICc as a sensitivity check and cite the matrix-pencil and Prony baselines.[4], [5]

VII. ABLATIONS

We ablate (i) heuristic vs MLP-only vs auto arbitration, (ii) RK4 vs Euler integration step size in propagation, and (iii) BIC vs fixed- K in ringdown. Auto mode offers the best latency-accuracy tradeoff; RK4 stabilizes bounce counts at coarse steps; BIC avoids overfitting short bursts.

For the anomaly detector, auto-arbitration improved throughput by up to 40% compared to MLP-only processing while maintaining equivalent accuracy. The performance gain was most pronounced under high load conditions when the probability values were far from the decision threshold, allowing the system to bypass the more expensive MLP computation.

In the ray tracing component, RK4 integration required approximately 20% more computation time per step compared to Euler integration, but allowed for 3-4x larger step sizes while maintaining the same propagation path accuracy. This

²We enforce a minimum frequency separation and discard roots outside the stability wedge.

TABLE II: Operational snapshot under steady load.

Metric	Value	Note
Cache hit ($M(z)$)	82%	15 min TTL
Max RPS @ $p95 \leq 6.4$ ms	920	HPA=8 pods
p99 latency	8.9 ms	/v1/score
CPU / pod	0.64	request/limit 1/2
Mem / pod	410 MiB	model shared




Fig. 1: Cache efficiency vs. p95 latency. Points labeled by pod count; dashed line is the p95 SLO.

resulted in a net efficiency gain of approximately 2.5x for typical atmospheric profiles.

VIII. OPERATIONAL NOTES: SCALING AND SLOS

IX. OPERATIONAL NOTES: SCALING AND SLOS

Process model. We deploy FastAPI+Uvicorn with `-workers=2k` where $k=\text{num_physical_cores}$; each worker pins NumPy/SciPy BLAS threads to 1. Read-only models live in a shared memory segment to minimize cold-start.

Kubernetes. We use requests/limits of `cpu: 1/2, memory: 512Mi/1Gi` per pod; HPA: target 70% CPU, min/max replicas [2, 20]. Liveness/readiness probes at `/v1/healthz`. Timeouts: 50 ms server read, 100 ms overall for `/v1/score`, 150 ms for `/v1/modes`, 200 ms for `/v1/propagate`. We cap `max_distance_m` and enforce per-request sampling budgets to hit the p95 6.4 ms SLO at 820 RPS.

```
# HPA tuned for p95<=6.4ms while keeping cost low
kubectl autoscale deploy rf-ghost-api \
  --cpu-percent=70 --min=2 --max=20
```

Caching. A 15 min TTL cache for nearest reference atmosphere or last radiosonde sounding avoids recomputation of $M(z)$; input-equivalent requests are deduplicated by a normalized hash of $\{z\text{-grid}, M(z), \text{step}\}$.

Failure policy. On `/v1/modes`, if BIC selects $k > 3$ at $\text{SNR} < 6$ dB, we fall back to $k \in \{1, 2\}$ with AICc and increase the minimum separation constraint.

X. RELATED WORK

Prior work covers efficient attention for spectra, classical propagation models, and parametric transient analysis. Our aim




Fig. 2: **Service diagram.** Solid arrows: request flow; dashed: cached context. All angles in degrees, heights in meters, M in M-units, time in milliseconds.

is the glue: a portable API that fuses score, propagation, and ringdown under tight p95 budgets.

Recent advances in ML-augmented propagation modeling [6] have demonstrated impressive accuracy for ducting prediction but typically require significant computational resources. Our approach strikes a balance between physical modeling and operational requirements, providing sufficient accuracy for anomaly contextualization without excessive latency.

In the domain of signal characterization, previous work has explored various approaches to transient signal decomposition, including wavelet analysis, empirical mode decomposition, and parametric methods. Our ringdown mode fitting approach builds on these foundations while optimizing for the specific characteristics of RF burst anomalies.

XI. CONCLUSION

Adding propagation and ringdown context to an anomaly service improves operator trust without sacrificing latency. Future work: learned priors from weather feeds and joint training of the ringdown head.

Our implementation demonstrates that modern microservice architectures can effectively integrate physical modeling with machine learning approaches to create more robust and interpretable anomaly detection systems. The clear API boundaries and modular design facilitate ongoing development and integration with existing operational workflows.

A promising direction for future research is the incorporation of real-time atmospheric data from weather services to dynamically update the propagation models. Additionally, end-to-end training of neural components that directly incorporate physical constraints could further improve system performance while maintaining interpretability.




Fig. 3: **Horizontal layout.** Shows the connections between client applications, API endpoints, and the various backend components with precise data flow patterns and processing units.

TABLE III: Endpoint return fields (units). Arrays shown with [].

Endpoint	Return Field	Description
/v1/score	p	Calibrated anomaly prob. (0–1)
	decision	Binary decision (0,1)
	latency_ms	Processing time (ms)
/v1/propagate	mode	Backend used (heuristic, mlp)
	path[]	Path points (x, z, m , bounce)
	ducted	Atmospheric ducting flag (bool)
	bounces[]	Bounce locations ($[x, z]$)
/v1/modes	max_range_m	Maximum propagation (m)
	confidence	Ducting confidence (0–1)
	freq_hz	Mode frequency (Hz)
	tau_s	Decay time constant (s)
	amp_rel	Relative amplitude (norm.)
	phase_rad	Phase offset (rad)

TABLE IV: Service SLO snapshot (auto-filled).

Metric	Value	Note
p50 latency	2.9 ms	Median request
p95 latency	6.4 ms	SLO threshold
Throughput	820 RPS	Sustained at SLO
Concurrency	24	Workers / inflight
Cold start	140 ms	First-hit warmup

APPENDIX A: JSON SCHEMAS

/v1/score → 200 OK

```
{
  "results": [
    {
      "p": float, "decision": 0|1, "latency_ms": float, "mode": "heuristic"|"mlp"
    }
  ],
  "p50_ms": float, "p95_ms": float, "backend": "heuristic"|"mlp"
}
```

/v1/propagate → 200 OK

```
{
  "path": [{"x_m": float, "z_m": float, "theta_rad": float, "m": float, "bounce": bool}],
  "ducted": bool, "inversion_detected": bool,
  "bounces": [[{x_m, z_m}, ...],
  "max_range_m": float, "confidence": float
}
```

```
"bounces": [[x_m, z_m], ...],
"max_range_m": float, "confidence": float
}
```

/v1/modes → 200 OK

```
{
  "modes": [{"freq_hz": float, "tau_s": float, "amp_rel": float, "phase_rad": float}]
}
```

Smoke tests (curl).

```
curl -s localhost:8080/v1/healthz
```

```
curl -s -X POST localhost:8080/v1/propagate \
-H 'content-type: application/json' \
-d '{ "azimuth_deg": 0, "elevation_deg": 1.0, "tx_pos": { "x_m": 100000, "y_m": 100000, "z_m": 100000 }, "max_distance_m": 150000, "step_m": 500, "sounding": [{ "z_m": 0, "N": 315 }, { "z_m": 50, "N": 315 } ] }
```

```
curl -s -X POST localhost:8080/v1/modes \
-H 'content-type: application/json' \
-d '{ "fs_hz": 10000, "signal": [0.1, 0.2, 0.0, -0.1, ...] }
```

APPENDIX

We integrate $\dot{x} = \cos \theta$, $\dot{z} = \sin \theta$, $\dot{\theta} = \kappa(M, z)$ where κ depends on dM/dz . Euler uses one slope; RK4 uses four stages $k_1, \dots, k_4$ and the weighted sum $(k_1 + 2k_2 + 2k_3 + k_4)/6$, yielding $O(h^4)$ local truncation error and improved bounce-point stability at coarse h .

REFERENCES

- [1] "Recommendation ITU-R P.453-14: The radio refractive index: its formula and refractivity data," <https://www.itu.int/rec/R-REC-P453-14-201908-I/en>, 2019, in force.
- [2] "Recommendation ITU-R P.835-7: Reference standard atmospheres," <https://www.itu.int/dms/pubrec/P835-7-202408-I/TOC-HTML-E.htm>, 2024, in force.
- [3] Y. Liu, H. Zhang, and B. Yan, "Ducting criteria and modified refractivity gradients," <https://pmc.ncbi.nlm.nih.gov/articles/PMC10221489/>, 2023, see eqs. relating $dM/dz = dN/dz + 0.157$ and trapping condition.
- [4] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 38, no. 5, pp. 814–824, 1990.

TABLE V: Endpoint return fields (units). Arrays shown with [].

Endpoint	Return Field	Description	Usage
/v1/score	p	Calibrated anomaly prob. (0–1)	Alert triggering, anomaly detection
	decision	Binary decision (0,1)	Quick thresholding for systems
	latency_ms	Processing time (ms)	Performance monitoring
	mode	Backend used (<i>heuristic</i> , <i>mlp</i>)	Execution pathway tracking
/v1/propagate	path[]	Path points (x, z, m , bounce)	Visualization, range calculation
	ducted	Atmospheric ducting flag (bool)	Ghost candidate flagging
	bounces[]	Bounce locations ($[x, z]$)	Multi-path analysis, terrain interaction
	max_range_m	Maximum propagation (m)	Coverage mapping, horizon calculation
/v1/modes	confidence	Ducting confidence (0–1)	Reliability metric for ducting assessment
	freq_hz	Mode frequency (Hz)	Signal characterization, transmitter ID
	tau_s	Decay time constant (s)	Material/path distance estimation
	amp_rel	Relative amplitude (norm.)	Power distribution across modes
	phase_rad	Phase offset (rad)	Coherence analysis, demodulation

- [5] “Prony’s method: damped sinusoid decomposition,” https://en.wikipedia.org/wiki/Prony%27s_method, 2024, background and historical context.
- [6] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On calibration of modern neural networks,” in *ICML*, 2017.