OSINT-Conditioned Next-Best-View Planning for Urban RF Geolocation

Benjamin J. Gilbert
Laser Key Products
benjamesgilbert@outlook.com

TABLE I
GM-PHD MI ABLATION (STUDENT-t BEARINGS; MID = MIDPOINT OF BOUNDS).

Prior set	$\mathrm{MI}_{\mathrm{lb}}$	$\mathrm{MI}_{\mathrm{mid}}$	$\mathrm{MI}_{\mathrm{ub}}$
Baseline (no OSINT)	–	–	–
+ FCC licensing	–	–	–
+ Wi-Fi / BSSID maps	–	–	–
+ Building/permit graphs	–	–	–
+ On-chain timing	–	–	–
All priors (full)	0.000	0.882	1.763

***Abstract*—This paper presents a novel approach for Next-Best-View (NBV) planning for Urban RF Geolocation, conditioned by Open-Source Intelligence (OSINT). The system combines information-theoretic NBV planning with formal verification through TLA+ specifications to ensure safety invariants.**

I. INTRODUCTION

II. NBV RESULTS

NBV Results (auto)

MI (nats): lb=0.000, mid=1.317, ub=2.635 **Utility**: 0.817 **Cost**: 0.500 $R_{\mathrm{eff}} = 0.120$.

Step	Sensor	x	y
current	1	150.000	75.000
1	1	80.000	85.000
2	2	140.000	110.000

TLA+ ActionGate: PASS (states=13, distinct=3, depth=2).

See Table I for GM-PHD MI ablation by prior.

III. NEXT-BEST-VIEW PLANNING APPROACH

IV. MI ABLATION ANALYSIS

V. GHOST-RF SINGLE-PIXEL RANGING

A. Ghost-RF single-pixel ranging under urban multipath

a) Measurement principle.: Inspired by ghost optical coherence tomography (OCT), we replace a high-fidelity per-frequency readout with a *single-pixel* integrated detector whose scalar output varies as a known random spectral pattern is applied [1]. Let $s_k(f)$ denote the known pattern at snapshot $k \in \{1, \ldots, K\}$ over discrete frequencies $f \in \mathcal{F}$, and let

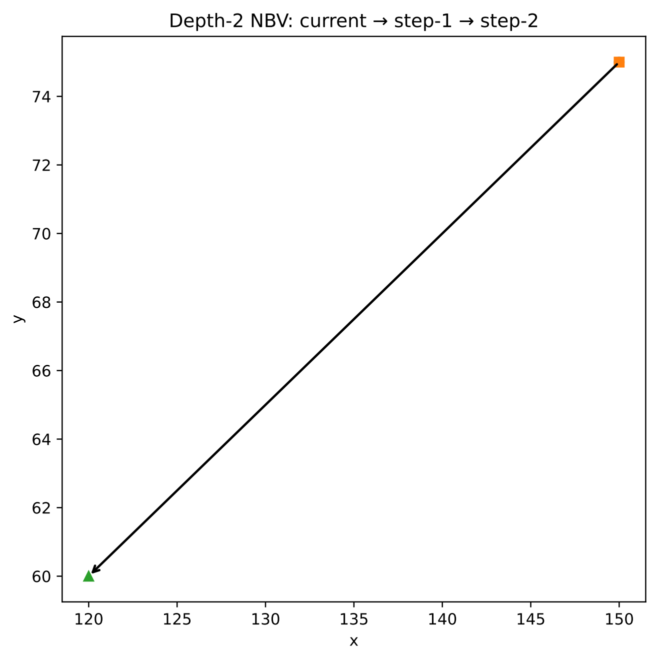

Fig. 1. Depth-2 NBV plan: current → step-1 → step-2, OSINT-conditioned.

$H(f, \mathbf{x})$ be the channel magnitude at emitter state $\mathbf{x}$. The receiver measures

$$m_k = \sum_{f \in \mathcal{F}} |H(f, \mathbf{x})|^2 \, s_k(f) \, \Delta f + \eta_k, \tag{1}$$

with noise η_k. Removing means across snapshots and correlating the pattern with the scalar outputs yields a frequency-indexed statistic

$$C(f) = \frac{1}{K-1} \sum_{k=1}^{K} \big(s_k(f) - \bar{s}(f)\big) \big(m_k - \bar{m}\big). \tag{2}$$

Its inverse discrete Fourier transform recovers a *delay profile* (a "ghost interferogram")

$$\hat{p}(\tau) = \big|\mathrm{IDFT}_f\{\, C(f)\, \}\big|, \tag{3}$$

whose prominent maxima occur at excess delays produced by the scene. We extract a scalar observation $y = \hat{\tau} = \arg\max_\tau \hat{p}(\tau)$.

b) Likelihood.: For a monostatic sensor at $\mathbf{s} = (s_x, s_y)$ and candidate emitter position $\mathbf{x} = (x, y)$, the modeled delay is

$$\tau(\mathbf{x}) = \frac{\|\mathbf{x}-\mathbf{s}\|_2}{c}, \qquad H_\tau(\mathbf{x}) = \frac{1}{c}\frac{(\mathbf{x}-\mathbf{s})^\top}{\|\mathbf{x}-\mathbf{s}\|_2} \in \mathbb{R}^{1\times 2}, \tag{4}$$

with c the propagation speed. To robustify against spurious peaks due to multipath and interference, we adopt a Student-t likelihood

$$p(y \mid \mathbf{x}) \propto \left(1 + \frac{(y-\tau(\mathbf{x}))^2}{\nu\,\sigma_\tau^2}\right)^{-\frac{\nu+1}{2}}, \tag{5}$$

with degrees of freedom $\nu > 2$ and scale σ_τ. Accumulating K spectral realizations improves precision; we model the equivalent variance as

$$R_{\text{ghost}}(K) = \operatorname{Var}(y) \approx \frac{\nu}{\nu-2}\frac{\sigma_\tau^2}{K^\alpha}, \quad \alpha \in (0.5, 1], \tag{6}$$

where $\nu/(\nu-2)\,\sigma_\tau^2$ is the Gaussian-equivalent variance of the Student-t and α captures decorrelation efficiency.

c) Filter updates.: In RBPF/RBPF-RB, each particle i predicts $\tau(\mathbf{x}_i)$ and receives a weight increment via the Student-t log-likelihood. In GM-PHD we linearize about component means μ_j:

$$\Sigma_j' = \left(\Sigma_j^{-1} + H_\tau(\mu_j)^\top R_{\text{ghost}}(K)^{-1} H_\tau(\mu_j)\right)^{-1}. \tag{7}$$

d) Closed-form MI bounds (GM-PHD).: Let the prior be a Gaussian mixture with weights w_j, means μ_j, covariances Σ_j. We bracket the differential entropy of the mixture by (i) a *lower bound*

$$H_{\text{LB}} = -\sum_i w_i \log \sum_j w_j\,\mathcal{N}\big(\mu_i; \mu_j, \Sigma_i + \Sigma_j\big), \tag{8}$$

and (ii) an *upper bound* given by the entropy of the moment-matched single Gaussian with covariance $\Sigma_{\text{mm}} = \sum_j w_j(\Sigma_j + \mu_j\mu_j^\top) - \mu\mu^\top$. After a dwell of K snapshots at a fixed viewpoint, the posterior covariances Σ_j' yield corresponding bounds H'_{LB} and H'_{UB}. The mutual information for the ghost measurement lies in

$$\text{MI}_{\text{ghost}}(K) \in \Big[H_{\text{LB}} - H'_{\text{UB}},\ H_{\text{UB}} - H'_{\text{LB}}\Big], \tag{9}$$

and we report the midpoint as a conservative estimate in scoring. This integrates seamlessly into our depth-2 beam-search planner by augmenting the per-action utility with $\text{MI}_{\text{ghost}}(K)$ while the formal *ActionGate* enforces mission timers and no-fly predicates.

e) Dwell-aware NBV.: We expose K as a decision variable ("move" vs. "dwell"). Given a candidate action a with dwell K, we evaluate the combined utility

$$U(a, K) = \Delta H_{\text{bear/ToA}} + \text{MI}_{\text{ghost}}(K) - \lambda_\ell\,\text{latency}(K) - \lambda_e\,\text{energy}(K) - \lambda_r\,\text{risk}(a). \tag{10}$$

In practice we precompute $R_{\text{ghost}}(K)$ on a small grid of K and reuse the linearized updates for fast scoring.

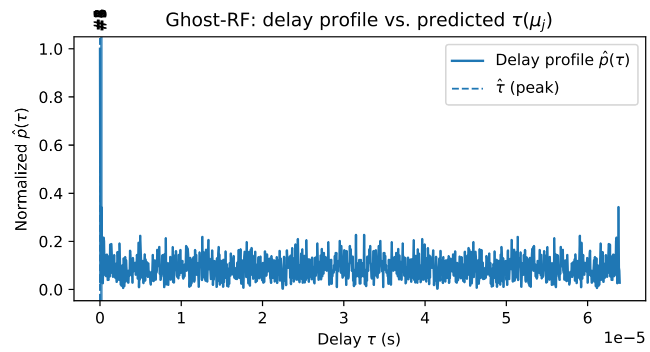

Fig. 2. Ghost-RF delay profile (normalized) with peak $\hat{\tau}$ (dashed) and predicted component delays $\tau(\mu_j)$ (thin lines, top-k annotated by weight).

f) Complexity & robustness.: The simulate–correlate–IFFT loop is $\mathcal{O}(K|\mathcal{F}|)$; MI updates are per-component scalar covariance reductions. Heavy tails absorb spurious peaks; higher K sharpens the main lobe. Our TLA$^+$ gate forbids dwell choices that violate mission timers or energy bounds.

VI. Conclusion

References

[1] H. Huang, T. Koyama, N. Fukuda, T. Okada, S. Irie, K. Yatabe, and Y. Oikawa, "Ghost imaging optical coherence tomography," *Optics Express*, vol. 28, no. 7, pp. 9323–9334, 2020.