Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Physics Informed Gaussian Process Radio Maps for Risk Aware UAV Trajectory Planning

G. Arul Dalton

aruldalton@saveetha.ac.in

Saveetha Engineering College

Research Article

Keywords: Radio map, Cellular-connected UAV, Gaussian process, Physics-informed learning, Uncertainty quantification, Trajectory planning.

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11480747/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Physics Informed Gaussian Process Radio Maps for Risk Aware UAV Trajectory Planning

G. Arul Dalton[1]*

[1] Department of Computer Science and Engineering, Saveetha Engineering College, Chennai, Tamil Nadu, India. E-mail: aruldalton@saveetha.ac.in

***Corresponding author:** G. Arul Dalton, e-mail: aruldalton@saveetha.ac.in

ORCID: G. Arul Dalton, 0000-0002-5981-7176

Abstract

Cellular-connected unmanned aerial vehicles (UAVs) flying at low altitude experience line-of-sight (LoS) interference from many cells, altitude-dependent shadowing and reception through the side lobes of down-tilted base station antennas, so safe operation requires an accurate three-dimensional (3D) radio map built from few measurements. This paper proposes SKYMAP, a framework that couples an antenna-aware, LoS-conditioned path-loss prior with a multicell Gaussian process residual field whose kernel encodes inter-site shadowing correlation and LoS- and altitude-dependent shadowing variance. The posterior map is converted into per-location probabilities that the signal-to-interference-plus-noise ratio falls below a threshold, and these outage probabilities drive two planners: a risk-weighted shortest-path planner and a new length-budgeted risk-aware planner (LB-RAP) that uses bisection to find the most outage-averse route whose length stays within a prescribed budget. In simulations over five statistically generated cities with the Third Generation Partnership Project aerial channel model, SKYMAP reduces the root-mean-square map error by 10.5 to 13.1% relative to physics-residual kriging and by 48 to 57% relative to ordinary kriging at sampling ratios of 0.5 to 5%. Its learned hyper-parameters recover the simulated channel statistics, and its predictive uncertainty is well calibrated. Risk-aware planning lowers the fraction of flight in outage from 37.3% on the shortest path to 2.5%, while LB-RAP met every length budget and reached 5.4% outage with only 7.6% extra path length. Limitations under model mismatch and for adaptive sampling are also reported.

Keywords

Radio map, Cellular-connected UAV, Gaussian process, Physics-informed learning, Uncertainty quantification, Trajectory planning.

1 Introduction

Unmanned aerial vehicles (UAVs) are becoming personal and professional communication devices in their own right. Delivery drones, inspection platforms and aerial cameras all need a reliable command-and-control link and a broadband data link while they fly in the low-altitude airspace, roughly 30 m to a few hundred metres above ground. Serving these vehicles through existing cellular networks is attractive because the infrastructure is already deployed, but terrestrial base stations (BSs) were designed for ground users [1–3]. A UAV above the rooftops sees many BSs in line of sight (LoS), which raises both the desired signal and the interference. It is also served mainly by the side lobes of antennas that are mechanically down-tilted toward the ground [3, 4]. As a result, the signal-to-interference-plus-noise ratio (SINR) of a UAV often decreases with altitude and changes abruptly as buildings create or remove LoS paths. Anticipating such link degradation is also central to mobility management in 5G networks, where learning-based models are used to make proactive handover decisions [5].

A radio map, also called a channel knowledge map, records location-specific channel quality. It allows the network or the UAV to anticipate coverage holes before reaching them [6, 7]. For UAVs the map

must be three-dimensional. It must also be built from sparse measurements, because every sample costs flight time, and it must state how confident it is, because a confident but wrong map leads directly to link loss. Existing radio-map estimators address only subsets of these needs. Interpolation methods such as inverse-distance weighting and kriging make weak assumptions but ignore propagation physics [8]. Deep models such as RadioUNet [9] are accurate in two dimensions but need large training corpora of simulated maps and give no calibrated uncertainty. Segmented regression [10] and radio-map-based path planning [11, 12] exploit LoS structure but rely on deterministic maps. Learning-based UAV trajectory planners have also been developed for obstacle avoidance and energy awareness [13], but they do not model cellular connectivity. Neural radio-frequency fields inspired by NeRF [14, 15] represent the channel as a coordinate network, but they are typically trained from dense measurements.

This paper asks a practical question. When only a few hundred UAV measurements are available, what is the most accurate and most trustworthy way to build a 3D radio map, and how should a UAV use the uncertainty of that map when it plans a route under a limited flight-length budget? We answer with SKYMAP (**S**parse s**K**y radio **MAP**ping). Its design principle is to put known physics into the mean, put the statistics of what is unknown into the covariance, and let the planner reason about risk rather than about point estimates.

The main contributions of this paper are as follows.

- **Antenna-aware, LoS-conditioned physics prior.** The mean function combines the known BS vertical antenna pattern with a LoS-conditioned two-slope path-loss model evaluated on the 3D city geometry. Its four path-loss coefficients are calibrated per cell by ridge-regularised least squares, so no site-specific tuning is required.
- **Multicell, LoS-heteroscedastic Gaussian process residual field.** The residual is modelled jointly across cells with a separable kernel that encodes inter-site shadowing correlation, a shadowing standard deviation that depends on the LoS state and decays with altitude as in 3GPP TR 36.777 [4], and anisotropic correlation lengths. All seven hyper-parameters are learned by type II maximum likelihood and have direct physical meaning.
- **Two uncertainty-aware planners.** The posterior map is converted into per-location outage probabilities by Monte Carlo sampling of the best-server SINR. A risk-weighted planner uses these probabilities as edge costs, and a new length-budgeted risk-aware planner (LB-RAP, Algorithm 3) finds the most outage-averse route that respects a hard path-length budget. LB-RAP is supported by a monotonicity proposition.
- **Transparent evaluation.** We compare against nine map estimators, including two neural radio fields, over five independent cities, with ablations, calibration analysis, a model-mismatch stress test, an adaptive-sampling study and 200 planning queries. We report negative results alongside positive ones, and we release all code.

The rest of the paper is organised as follows. Section 2 reviews related work. Section 3 presents the system model. Section 4 develops SKYMAP and the two planners. Section 5 describes the experimental setup, Section 6 reports the results, Section 7 discusses limitations, and Section 8 concludes.

2 Related Work

2.1 Radio Map Estimation

Radio map estimation, also known as spectrum cartography, reconstructs a spatial field of received power or channel gain from scattered measurements [6, 8]. Classical approaches include inverse-distance weighting, kriging, and matrix or tensor completion. Kriging is the geostatistical best linear unbiased predictor under a stationary covariance model [16] and is equivalent to Gaussian process regression [17]. Regression kriging adds a deterministic trend, typically a log-distance path-loss model, before kriging the residuals. We use this physics-residual kriging (PR-Kriging) as our strongest baseline. Levie et al. [9] proposed RadioUNet, a convolutional network that maps city layouts to two-dimensional path-loss maps, and Zeng et al. [7] surveyed channel knowledge maps as an enabler of environment-aware 6G.

2.2 Radio Maps and Trajectories for UAVs

Chen et al. [10] learned UAV radio maps by segmented regression over LoS and non-line-of-sight (NLoS) classes. Zhang and Zhang [11] used a radio map for 3D path planning of a cellular-connected UAV, minimising flight time subject to an outage constraint. Zeng et al. [12] proposed simultaneous navigation and radio mapping with deep reinforcement learning. Archana et al. [13] studied energy-aware adaptive obstacle avoidance for UAV trajectory planning with meta-reinforcement learning and segmentation. Manju et al. [18] addressed decentralized control design for communication in UAV swarms. These works establish the value of maps, learning and coordinated control for UAV navigation and communication, but they use deterministic map estimates and do not price map uncertainty in the planning cost.

2.3 Neural and Physics-Informed Radio Fields

Coordinate networks with Fourier features [19] underlie neural radiance fields [14] and their radio-frequency extension NeRF^2 [15]. Physics-informed neural networks [20] add soft constraints derived from governing equations. Both families are expressive, but their data requirements and uncertainty calibration in the sparse regime are not well understood. Uncertainty for neural models is usually obtained with deep ensembles [21]. We therefore include a Fourier-feature neural field and a physics-informed neural field (PINF) with the same prior and LoS features as SKYMAP among our baselines.

2.4 Research Gap

Table 1 summarises representative approaches. To the best of our knowledge, no existing UAV radio-map method jointly (a) builds a full 3D map, (b) embeds the antenna pattern and LoS geometry in the

prior, (c) shares statistical strength across cells through inter-site shadowing correlation, (d) provides calibrated uncertainty, and (e) uses that uncertainty in trajectory planning under an explicit length budget.

Table 1 Comparison of representative approaches (✓ supported, ✗ not supported, ~ partially supported)

Method	3D	Physics prior	Multicell	Calibrated uncertainty	Uncertainty-aware planning
Ordinary kriging [16]	✓	✗	✗	✓	✗
RadioUNet [9]	✗	~ (city map)	✗	✗	✗
Segmented regression [10]	✓	✓ (LoS)	✗	✗	✗
Radio-map path planning [11]	✓	✓	✗	✗	~ (deterministic map)
SNARM [12]	✓	✗	✗	✗	~ (learned policy)
Meta-RL obstacle avoidance [13]	✓	✗	✗	✗	✗ (no radio model)
NeRF² [15]	✓	✗	✗	✗	✗
PR-Kriging (baseline)	✓	✓	✗	✓	✗
SKYMAP (this work)	✓	✓ (antenna + LoS)	✓	✓	✓ (with length budget)

3 System Model and Problem Formulation

3.1 Network and Airspace

We consider a set $\mathcal{B} = \{1, \ldots, B\}$ of terrestrial BSs serving a low-altitude airspace $\mathcal{V} \subset \mathbb{R}^3$ above an urban area whose building footprints and heights are known from a 3D city model (Fig. 1). For a location $\mathbf{x} = (x, y, z) \in \mathcal{V}$ and BS b, let $d_b(\mathbf{x})$ denote the 3D distance and $\ell_b(\mathbf{x}) \in \{0, 1\}$ the geometric LoS indicator. The indicator is obtained by ray marching over the building height raster rather than from a statistical LoS-probability model [22]. Positions inside buildings are excluded from $\mathcal{V}$.

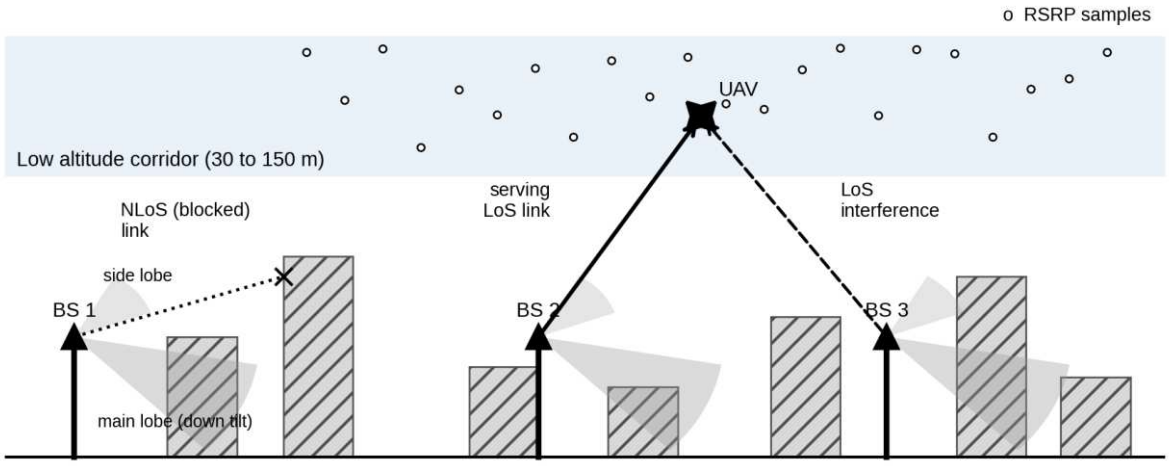

Fig. 1 System model (side view). A UAV in the low-altitude corridor is served by a LoS link through the side lobe of a down-tilted BS antenna, receives LoS interference from other cells, and may be blocked from some BSs by buildings. Open circles mark the sparse received-power measurement locations

3.2 Channel Model

The average received power (RSRP, in dBm) from BS b at $\mathbf{x}$ is

$$P_b(x) = P_T + G_b(x) - PL_b(x) + S_b(x), \quad (1)$$

where P_T is the transmit power, PL_b follows the LoS/NLoS aerial path-loss model of 3GPP TR 36.777 [4], S_b is log-normal shadowing, and G_b is the BS vertical antenna gain of 3GPP TR 38.901 [23]:

$$G_b(x) = G_{max} - min\{ 12\,[(\theta_b(x) - 90° - \theta_{tilt}) / \theta_{3dB}]^2, SLA \}, \quad (2)$$

where θ_b is the zenith angle from BS b to $\mathbf{x}$, θ_{tilt} is the electrical down-tilt, θ_{3dB} is the vertical half-power beamwidth and SLA is the side-lobe attenuation. Following [4], the shadowing standard deviation is $4.64 \cdot \exp(-0.0066\, z)$ dB in LoS and 6 dB in NLoS. Shadowing is spatially correlated [24] and partially correlated across sites [23]. A UAV at $\mathbf{x}$ attaches to the strongest BS b^*, and its SINR is

$$\gamma(x) = p_{b^*}(x) / (\Sigma_{b' \neq b^*} p_{b'}(x) + N_0), \quad (3)$$

where $p_b = 10^{\wedge}(P_b/10)$ is in milliwatts and N_0 is the receiver noise power. Location $\mathbf{x}$ is in outage if $\gamma(\mathbf{x}) < \gamma_{th}$.

3.3 Measurement Model

A measurement campaign collects n samples $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}$, $i = 1, \ldots, n$. Each $\mathbf{y}_i \in \mathbb{R}^B$ contains the RSRP of all B cells at $\mathbf{x}_i$, as reported by standard user-equipment measurement procedures:

$$y_{i,b} = P_b(x_i) + \varepsilon_{i,b}, \quad \varepsilon_{i,b} \sim \mathcal{N}(0, \sigma_n^2). \quad (4)$$

3.4 Problem Statement

Given $\mathcal{D}$ and the city model, the objectives are as follows. (P1) Estimate the posterior distribution of $\{P_b(\mathbf{x})\}$ for every $\mathbf{x} \in \mathcal{V}$ with small reconstruction error and calibrated uncertainty. (P2) Given a start, a goal and a length budget δ, find a trajectory that minimises the expected fraction of flight in outage while its length stays within $(1 + \delta)$ times the shortest feasible path. "Sparse" means $n/|\mathcal{V}|$ between 0.5% and 5%.

4 Proposed SKYMAP Framework

4.1 Overview

Fig. 2 shows the two stages of SKYMAP. In Stage I (Algorithm 1), the measurements, the city model and the BS antenna configuration feed a physics prior. A multicell Gaussian process (GP) then models what the prior cannot explain, and the result is a posterior mean and standard deviation for every location and cell. In Stage II, the posterior is converted into outage probabilities, which define edge costs on a 3D graph. The risk-weighted planner (Algorithm 2) or the length-budgeted planner LB-RAP (Algorithm 3) then computes the route.

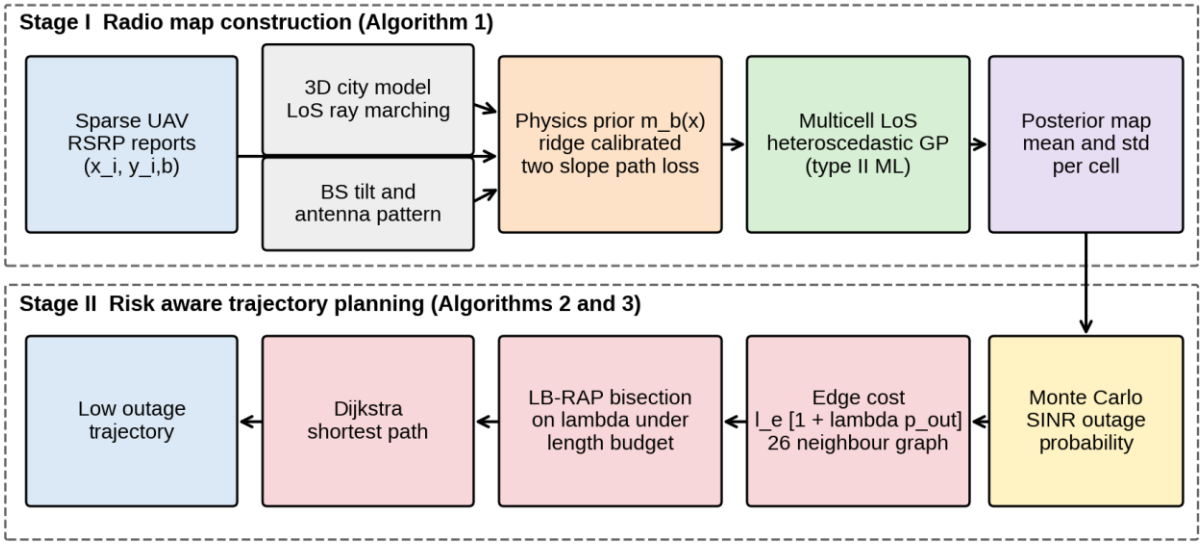

Fig. 2 Block diagram of the proposed SKYMAP framework. Stage I constructs the probabilistic 3D radio map (Algorithm 1), and Stage II converts it into outage probabilities and plans risk-aware trajectories with the risk-weighted planner (Algorithm 2) or the length-budgeted planner LB-RAP (Algorithm 3)

4.2 Antenna-Aware, LoS-Conditioned Physics Prior

The mean function for cell b is

$$m_b(x) = G_b(x) + \ell_b(x)[\alpha_{L,b} + \beta_{L,b}\cdot 10\log_{10} d_b(x)] + (1 - \ell_b(x))[\alpha_{N,b} + \beta_{N,b}\cdot 10\log_{10} d_b(x)]. \quad (5)$$

The antenna term G_b uses only the operator-known tilt and beamwidth, and the LoS indicator uses only the city model. The coefficients $\boldsymbol{\theta}_b = [\alpha_{L,b}, \beta_{L,b}, \alpha_{N,b}, \beta_{N,b}]^T$ absorb the transmit power, frequency-dependent loss and environment-specific slopes. They are estimated by ridge regression toward a pooled log-distance fit $\bar{\boldsymbol{\theta}}_b$, which stabilises the estimate when one LoS class is sparsely sampled:

$$\hat{\theta}_b = \operatorname{argmin}_\theta \|y_b - g_b - A_b\,\theta\|^2 + \eta\|\theta - \bar{\theta}_b\|^2, \quad (6)$$

where $\mathbf{A}_b$ stacks the regressors of (5) at the measured locations and $\mathbf{g}_b$ the antenna gains. The minimiser has the closed form $\hat{\boldsymbol{\theta}}_b = (\mathbf{A}_b^T\mathbf{A}_b + \eta\mathbf{I})^{-1}(\mathbf{A}_b^T(\mathbf{y}_b - \mathbf{g}_b) + \eta\bar{\boldsymbol{\theta}}_b)$, and we use $\eta = 1$.

4.3 Multicell, LoS-Heteroscedastic Residual Field

Let $r_b(\mathbf{x}) = P_b(\mathbf{x}) - m_b(\mathbf{x})$. We place a zero-mean joint GP prior on $\{r_b\}$ with covariance

$$\mathrm{Cov}[r_b(x), r_{b'}(x')] = s_b(x)\, s_{b'}(x')\,[(1-\rho)\,\delta_{bb'} + \rho\,]\,k(x, x'), \quad (7)$$

which is a linear model of coregionalisation [25] with an equicorrelated coregionalisation matrix. The parameter ρ captures inter-site shadowing correlation, which 3GPP models as 0.5 [23]. The LoS-heteroscedastic scale mirrors the 3GPP altitude dependence:

$$s_b(x) = \sigma_L\cdot\exp(-\kappa z / 100) \text{ if } \ell_b(x) = 1, \quad s_b(x) = \sigma_N \text{ otherwise.} \quad (8)$$

The spatial correlation uses an anisotropic Matérn-3/2 kernel, which is rougher than the squared-exponential kernel and therefore better suited to shadowing fields:

$$k(x, x') = (1 + \sqrt{3}\,r)\exp(-\sqrt{3}\,r), \quad r^2 = [(\Delta x)^2 + (\Delta y)^2]/\ell_h^2 + (\Delta z)^2/\ell_v^2. \quad (9)$$

The heteroscedastic scales multiply a positive-definite kernel, and the coregionalisation matrix is positive definite for $\rho \in (-1/(B-1), 1)$. Hence (7) is a valid covariance.

4.4 Hyper-parameter Learning

The hyper-parameters $\psi = \{\ell_h, \ell_v, \rho, \sigma_L, \kappa, \sigma_N, \sigma_n\}$ are learned by maximising the log marginal likelihood of the stacked residual vector $\mathbf{r} \in \mathbb{R}^{nB}$ [17]:

$$\log p(r \mid \psi) = -\tfrac{1}{2}\, r^T (K_\psi + \sigma_n^2 I)^{-1} r - \tfrac{1}{2} \log |K_\psi + \sigma_n^2 I| - (nB/2) \log 2\pi. \tag{10}$$

We optimise (10) with box-constrained L-BFGS-B on a random subset of at most $n_h = 300$ locations (900 observations), which keeps each likelihood evaluation cheap, and then condition on all n samples.

4.5 Posterior Radio Map

For a query location $\mathbf{x}$ and cell b, the posterior mean and variance are

$$\mu_b(x) = m_b(x) + k_b(x)^T (K + \sigma_n^2 I)^{-1} r, \tag{11}$$

$$\sigma_b^2(x) = s_b(x)^2 - k_b(x)^T (K + \sigma_n^2 I)^{-1} k_b(x), \tag{12}$$

where $\mathbf{k}_b(\mathbf{x}) \in \mathbb{R}^{nB}$ collects the cross-covariances (7) between $(\mathbf{x}, b)$ and all observations. A single Cholesky factorisation is shared by all queries. Algorithm 1 summarises Stage I.

Algorithm 1 SKYMAP radio-map construction

Input: measurements $\mathcal{D} = \{(x_i, y_i)\}$, city model, BS positions and tilts
Output: posterior mean $\mu_b(x)$ and standard deviation $\sigma_b(x)$ for all $x \in \mathcal{V}$, $b \in \mathcal{B}$
1: Compute LoS indicators $\ell_b(x)$ by ray marching and antenna gains $G_b(x)$ by (2)
2: **for** b = 1, …, B **do** calibrate $\hat{\theta}_b$ by ridge regression (6) **end for**
3: Form residuals $r_{i,b} = y_{i,b} - m_b(x_i)$ using (5)
4: Draw a subset of n_h locations; $\hat{\psi} \leftarrow$ L-BFGS-B maximisation of (10)
5: Build K over all nB observations; Cholesky $L\,L^T = K + \sigma_n^2 I$; $\alpha = L^{-T} L^{-1} r$
6: **for** each query batch X_q and cell b **do**
$\mu_b \leftarrow m_b(X_q) + K_{qo}\,\alpha$; $V \leftarrow L^{-1} K_{oq}$; $\sigma_b^2 \leftarrow s_b(X_q)^2 - \mathrm{colsum}(V \circ V)$
7: **end for**; **return** $\{\mu_b, \sigma_b\}$

4.6 Outage Probability and Risk-Weighted Planning

The planner needs the probability that a location is in outage, not merely a point estimate of SINR. Because (3) is a nonlinear function of the B received powers, we estimate it by Monte Carlo sampling from the posterior marginals:

$$p_{out}(x) = Pr[\gamma(x) < \gamma_{th} \mid \mathcal{D}] \approx (1/M) \Sigma_{m=1}\hat{}\{M\}\, \mathbb{1}[\gamma^{(m)}(x) < \gamma_{th}], \quad P_b^{(m)}(x) \sim \mathcal{N}(\mu_b(x), \sigma_b^2(x)). \tag{13}$$

The airspace is discretised into a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with a 26-neighbourhood. Every edge $e = (u, v)$ of length ℓ_e receives the cost

$$c_e(\lambda) = \ell_e \left[1 + \lambda \left(p_{out}(u) + p_{out}(v) \right) / 2 \right], \tag{14}$$

where $\lambda \geq 0$ weights outage exposure against path length. With $\lambda = 0$ the planner returns the shortest path. Replacing p_{out} with the hard indicator $\mathbb{1}[\hat{\gamma} < \gamma_{th}]$ of the posterior mean gives a deterministic mean-

map planner. All costs are non-negative, so Dijkstra's algorithm [26] returns the optimal path (Algorithm 2).

Algorithm 2 Risk-weighted trajectory planning

Input: posterior $\{\mu_b, \sigma_b\}$ on graph nodes $\mathcal{N}$, threshold γ_{th}, weight λ, samples M, start s, goal g
1: **for** each node $u \in \mathcal{N}$ **do** draw M samples of $\{P_b(u)\}$, compute $\gamma^{(m)}(u)$ by (3), and set $p_{out}(u)$ by (13)
2: **for** each edge $e \in \mathcal{E}$ **do** set $c_e(\lambda)$ by (14)
3: $\pi^* \leftarrow$ Dijkstra($\mathcal{G}$, c(λ), s, g); **return** π^*

4.7 Length-Budgeted Risk-Aware Planning (LB-RAP)

A fixed λ gives no control over how much longer the route becomes, yet flight time and battery energy are hard constraints for a UAV operator. We therefore pose the budgeted problem

$$\min_{\pi} R(\pi) \quad \text{subject to} \quad L(\pi) \le (1+\delta)\, L(\pi_0), \tag{15}$$

where $L(\pi) = \Sigma_{e\in\pi}\, \ell_e$ is the path length, $R(\pi) = \Sigma_{e\in\pi}\, \ell_e\, (p_{out}(u) + p_{out}(v))/2$ is the length-weighted outage exposure, and π_0 is the shortest path. The cost (14) is the Lagrangian $L(\pi) + \lambda_R(\pi)$ of (15), and the following property makes a simple search over λ possible.

Proposition 1 *Let π_1 and π_2 minimise $L + \lambda_1 R$ and $L + \lambda_2 R$, respectively, with $0 \le \lambda_1 < \lambda_2$. Then $R(\pi_1) \ge R(\pi_2)$ and $L(\pi_1) \le L(\pi_2)$.*

Proof. Optimality gives $L_1 + \lambda_1 R_1 \le L_2 + \lambda_1 R_2$ and $L_2 + \lambda_2 R_2 \le L_1 + \lambda_2 R_1$. Adding the two inequalities yields $(\lambda_1 - \lambda_2)(R_1 - R_2) \le 0$, hence $R_1 \ge R_2$. Substituting into the first inequality gives $L_1 \le L_2 + \lambda_1(R_2 - R_1) \le L_2$. ■

Path length is therefore non-decreasing and outage exposure non-increasing in λ, so the largest λ whose optimal path still satisfies the budget can be found by bisection (Algorithm 3). Each iteration is one Dijkstra call, so LB-RAP costs $(K + 2)$ shortest-path computations and always returns a feasible path, because π_0 itself is feasible.

Algorithm 3 LB-RAP: length-budgeted risk-aware planning (proposed)

Input: graph $\mathcal{G}$ with outage probabilities p_{out}, start s, goal g, budget δ, upper weight λ_{max}, iterations K
Output: path π^* with $L(\pi^*) \le (1+\delta)\, L(\pi_0)$ and minimal outage exposure among the Lagrangian solutions
1: $\pi_0 \leftarrow$ Algorithm 2 with $\lambda = 0$; $L_B \leftarrow (1+\delta)\, L(\pi_0)$
2: $\pi_{max} \leftarrow$ Algorithm 2 with $\lambda = \lambda_{max}$; **if** $L(\pi_{max}) \le L_B$ **then return** π_{max}
3: $\lambda_{lo} \leftarrow 0$; $\lambda_{hi} \leftarrow \lambda_{max}$; $\pi^* \leftarrow \pi_0$
4: **for** k = 1, …, K **do**
  $\lambda \leftarrow (\lambda_{lo} + \lambda_{hi})/2$; $\pi \leftarrow$ Algorithm 2 with weight λ
  if $L(\pi) \le L_B$ **then** $\lambda_{lo} \leftarrow \lambda$; $\pi^* \leftarrow \pi$ **else** $\lambda_{hi} \leftarrow \lambda$
5: **end for**; **return** π^*

4.8 Complexity

Each likelihood evaluation costs $O((n_{hB}_)^3)$. Conditioning costs $O((nB)^3)$ once, and each query costs $O(nB)$ for the mean and $O((nB)^2)$ for the variance. For $n = 1660$ and $B = 3$ (5% sampling), map

construction took 45 s on one CPU core (Section 6.1); sparse or inducing-point approximations [17] can reduce this for city-scale deployments. Computing (13) costs $O(MB|\mathcal{N}|)$ once. Algorithm 2 costs $O(|\mathcal{E}| + |\mathcal{N}| \log |\mathcal{N}|)$ per query, and Algorithm 3 costs $(K + 2)$ times that, which was below 0.11 s per query in our experiments.

5 Experimental Setup

5.1 Scenario Generation

We generated five independent urban environments (seeds 1 to 5) with the ITU-R P.1410 statistical model [27] ($\alpha = 0.3$, $\beta = 500$ buildings/km², $\gamma = 15$ m). A sixth, separate environment (seed 99) was used only to choose baseline hyper-parameters, so no test environment influenced any design choice. Table 2 lists the parameters. Ground-truth maps follow (1)–(3) with 3D-correlated shadowing and an inter-site correlation of 0.5, and measurements add 2 dB Gaussian noise. Sampling locations are drawn uniformly at random from the valid airspace, and accuracy is evaluated on all remaining valid grid points (about 31,500 to 33,000 points per cell).

Table 2 Simulation parameters

Parameter	Value
Area and map grid	600 m × 600 m; 10 m horizontal grid
UAV altitudes	30–150 m in 15 m steps (9 layers); 33,489 grid points
City model	ITU-R P.1410: $\alpha = 0.3$, $\beta = 500$ km^{-2}, $\gamma = 15$ m (heights 6–70 m)
Base stations	$B = 3$, rooftop, height ≥ 25 m, $P_T = 46$ dBm, $f_c = 2$ GHz
Antenna (TR 38.901)	$G_{max} = 8$ dBi, $\theta_{tilt} = 12°$, $\theta_{3dB} = 65°$, SLA = 30 dB
Path loss and shadowing	3GPP TR 36.777 UMa-AV; $\sigma_{LoS} = 4.64\ e^{\wedge}(-0.0066\ z)$ dB, $\sigma_{NLoS} = 6$ dB
Shadowing correlation	3D Gaussian-filtered field (about 30 m horizontal, 18 m vertical); inter-site $\rho = 0.5$
Noise	$N_0 = -97$ dBm (10 MHz, noise figure 7 dB); measurement noise $\sigma_n = 2$ dB
Sampling ratios	0.5%, 1%, 2%, 3%, 5% (166–1660 locations)
Planner	20 m graph, 26-neighbourhood, $\gamma_{th} = 0$ dB, $\lambda = 8$ (Algorithm 2), $M = 256$
LB-RAP	$\delta \in \{5\%, 10\%, 20\%\}$, $\lambda_{max} = 64$, $K = 12$
Planning queries	40 per city (200 in total); horizontal separation ≥ 400 m; start and goal altitude ≥ 60 m
Hardware and software	Single CPU core; Python with NumPy, SciPy and scikit-learn

5.2 Baselines

We compare SKYMAP with eight map estimators: (i) **LDPL**, a per-cell log-distance fit; (ii) **Two-slope**, the LoS-conditioned prior of (5) without the antenna or residual terms; (iii) **IDW**, inverse-distance weighting with 8 neighbours and power 2; (iv) **Kriging**, a per-cell GP with an anisotropic radial-basis kernel and white noise; (v) **PR-Kriging**, the two-slope prior plus per-cell GP residuals; (vi) **MLP**, a ReLU network (64-64) on normalised coordinates; (vii) **FF-NF**, a Fourier-feature neural field (48 frequencies, tanh 64-64); and (viii) **PINF**, a physics-informed neural field that takes the same two-slope prior, LoS indicators and log-distance features as input and adds smoothness, residual-energy and monotone radial-decay penalties. The neural fields were trained with Adam for 1200 iterations with

cosine learning-rate decay. For planning we compare the **Shortest** path, a **PR-Kriging** mean-map planner, a **SKYMAP-mean** planner, the **SKYMAP-risk** planner (Algorithm 2), **LB-RAP** (Algorithm 3) on both SKYMAP and PR-Kriging maps, and an **Oracle** planner that knows the true outage map.

5.3 Metrics

For map accuracy we report the root-mean-square error (RMSE) and mean absolute error (MAE) in dB, averaged over cells and test points. For uncertainty we report the prediction-interval coverage probability (PICP) at nominal 68% and 95%, the mean 95% interval width (MPIW) and the Gaussian negative log-likelihood (NLL) on held-out noisy measurements. For planning we report the length-weighted fraction of each path whose true SINR is below γ_{th}, the path length, the extra length relative to the shortest path, and the budget-satisfaction rate. All results are mean ± standard deviation over the five test cities.

5.4 Implementation

All methods were implemented in Python. The complete code, fixed random seeds and raw result files are available in the public repository listed under Data Availability, and every table and figure can be regenerated with the provided scripts.

6 Results and Discussion

6.1 Radio Map Accuracy

Table 3 and Fig. 3 report RMSE against sampling ratio. SKYMAP is the most accurate method at every ratio. It improves on PR-Kriging, the strongest baseline, by 12.7%, 13.1%, 10.5%, 10.5% and 11.8% at 0.5%, 1%, 2%, 3% and 5% sampling, respectively, and on ordinary kriging by 48 to 57%. SKYMAP beat PR-Kriging in all 25 paired (city, ratio) trials, and a Wilcoxon signed-rank test gives $p = 6.0 \times 10^{-8}$. With only 166 measurements (0.5%), SKYMAP (3.04 dB) is more accurate than ordinary kriging with ten times as many samples (3.27 dB at 5%).

Table 3 Radio-map RMSE (dB, mean ± standard deviation over five cities) versus sampling ratio; best values in bold

Method	0.5% (n = 166)	1% (n = 332)	2% (n = 664)	3% (n = 996)	5% (n = 1660)
LDPL	6.91 ± 0.52	6.90 ± 0.51	6.88 ± 0.50	6.87 ± 0.51	6.87 ± 0.51
Two-slope (LoS)	4.51 ± 0.31	4.44 ± 0.33	4.40 ± 0.31	4.39 ± 0.31	4.38 ± 0.31
IDW	5.51 ± 0.33	4.87 ± 0.30	4.22 ± 0.14	3.91 ± 0.17	3.47 ± 0.18
Kriging	5.83 ± 0.36	5.00 ± 0.28	4.15 ± 0.08	3.70 ± 0.13	3.27 ± 0.16
MLP	6.31 ± 0.35	5.22 ± 0.33	4.66 ± 0.24	4.44 ± 0.34	4.23 ± 0.32
FF-NF	6.94 ± 0.15	6.13 ± 0.23	5.70 ± 0.22	5.35 ± 0.18	4.52 ± 0.20
PINF	3.93 ± 0.25	3.39 ± 0.23	2.91 ± 0.22	2.66 ± 0.19	2.33 ± 0.18
PR-Kriging	3.49 ± 0.18	2.80 ± 0.21	2.22 ± 0.12	1.95 ± 0.09	1.60 ± 0.11
SKYMAP	**3.04 ± 0.22**	**2.44 ± 0.22**	**1.99 ± 0.11**	**1.75 ± 0.10**	**1.41 ± 0.10**

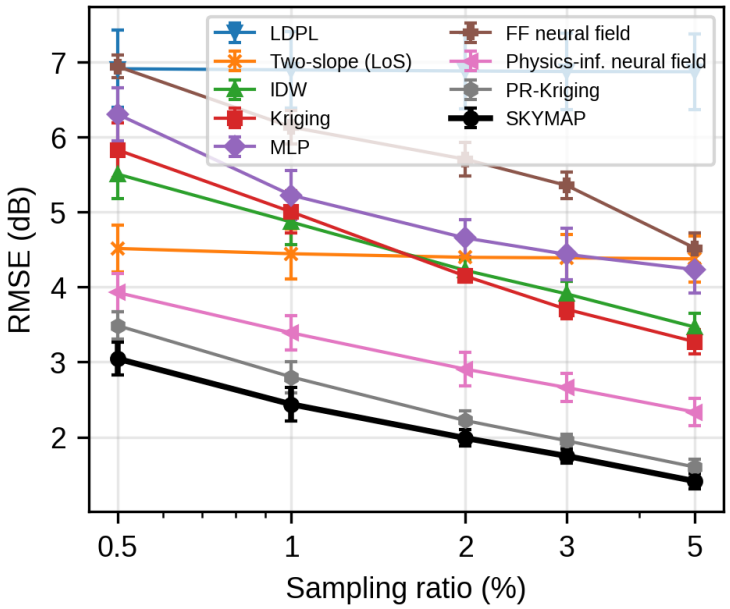

Fig. 3 RMSE versus sampling ratio on a logarithmic axis. Error bars show ±1 standard deviation over five cities, and SKYMAP is drawn with the thick line and circular markers

Three observations stand out. First, physics outweighs model capacity in the sparse regime. The unconstrained MLP and FF-NF are worse than the four-parameter two-slope model, and FF-NF overfits the measurement noise. Adding the physics prior and LoS features (PINF) closes most of this gap, but PINF still trails PR-Kriging by 0.5 to 0.7 dB, because a GP interpolates correlated shadowing more efficiently than a coordinate network trained by gradient descent. Second, the geometry-based prior (two-slope, 4.4 dB) is far better than the distance-only LDPL (6.9 dB), which confirms that the LoS state is the most important structural feature of the UAV channel. Third, the relative advantage of SKYMAP over PR-Kriging is roughly constant across sampling ratios, so it comes from better modelling rather than from a small-sample effect.

Fig. 4 shows reconstructions at 90 m altitude. SKYMAP reproduces the sharp LoS boundary and the shadowing pattern. Its posterior standard deviation is low near measurements and in LoS regions and high in NLoS pockets.

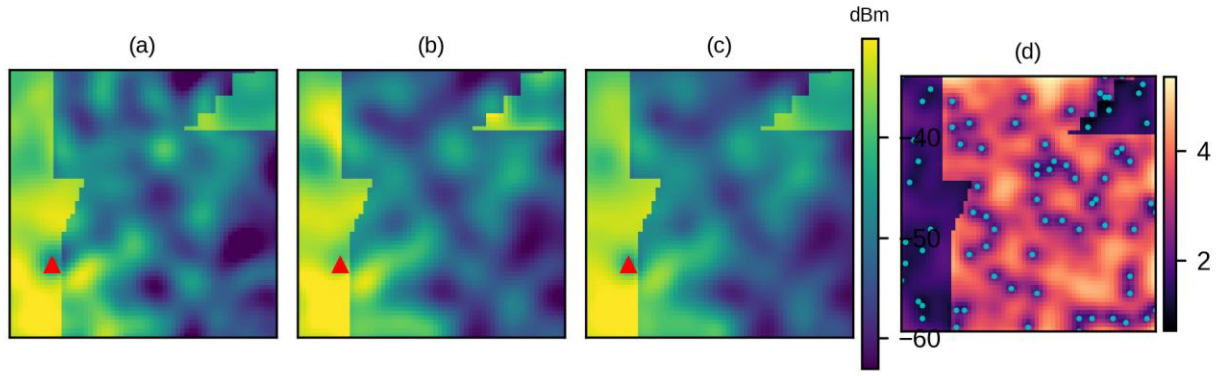

Fig. 4 RSRP (dBm) of cell 1 at 90 m altitude in city 1 with 2% sampling: (a) ground truth, (b) PR-Kriging, (c) SKYMAP posterior mean, and (d) SKYMAP posterior standard deviation in dB (dots mark the measurements in this layer; triangles mark the BS)

Table 4 lists MAE and run time. SKYMAP has the lowest MAE (1.44 dB versus 1.62 dB for PR-Kriging). Because its GP is joint across cells, it costs roughly 1.7 to 3 times the run time of PR-Kriging, but it still finishes within a minute on one CPU core at 5% sampling.

Table 4 MAE at 2% sampling and wall-clock time on a single CPU core

Method	MAE at 2% (dB)	Time at 2% (s)	Time at 5% (s)
IDW	2.60	0.04	0.04
Kriging	2.57	2.8	21.0
MLP	3.05	0.9	2.3
FF-NF	3.96	1.5	3.6
PINF	2.07	7.4	9.8
PR-Kriging	1.62	4.1	26.7
SKYMAP	**1.44**	12.1	44.8

6.2 Interpretability of Learned Hyper-parameters

Every SKYMAP hyper-parameter corresponds to a physical quantity, so we can check whether the model has learned the true channel statistics. Table 5 compares the values learned at 5% sampling with the simulator settings. The inter-site correlation ρ (0.45 versus 0.5), the LoS shadowing scale σ_L (4.79 versus 4.64 dB), its altitude decay κ (0.64 versus 0.66), the NLoS shadowing σ_N (6.24 versus 6.0 dB) and the measurement noise σ_n (1.85 versus 2.0 dB) are all recovered to within about one standard deviation. SKYMAP can therefore double as a measurement-driven estimator of channel parameters, which is useful for checking or recalibrating standardised aerial channel models in a specific city.

Table 5 Learned hyper-parameters (5% sampling, mean ± standard deviation over five cities) versus simulator ground truth

Hyper-parameter	Physical meaning	Learned	Ground truth
ρ	Inter-site shadowing correlation	0.45 ± 0.06	0.50
σ_L (dB)	LoS shadowing standard deviation at z = 0	4.79 ± 0.64	4.64
κ	LoS decay per 100 m of altitude	0.64 ± 0.12	0.66

Hyper-parameter	Physical meaning	Learned	Ground truth
σ_N (dB)	NLoS shadowing standard deviation	6.24 ± 0.39	6.00
σ_n (dB)	Measurement noise	1.85 ± 0.10	2.00
ℓ_h / ℓ_v (m)	Horizontal / vertical correlation length	50.9 ± 4.7 / 25.6 ± 2.8	not directly comparable (Gaussian field)

6.3 Uncertainty Calibration

Planning with risk requires trustworthy uncertainty. Table 6 and Fig. 5a evaluate the predictive distribution on held-out noisy measurements at 2% sampling. The 95% intervals of SKYMAP contain 95.0% of the observations and its 68% intervals contain 68.5%, which is essentially perfect calibration. Its intervals are also narrower and its NLL lower than those of PR-Kriging. The reason is the LoS-heteroscedastic scale, which assigns wide intervals to NLoS regions and narrow ones to high-altitude LoS regions instead of one average width everywhere.

Table 6 Predictive uncertainty at 2% sampling (mean over five cities)

Method	PICP 68% (ideal 0.683)	PICP 95% (ideal 0.950)	MPIW 95% (dB)	NLL
PR-Kriging	0.687	0.943	11.32	2.501
SKYMAP	0.685	**0.950**	**11.01**	**2.421**

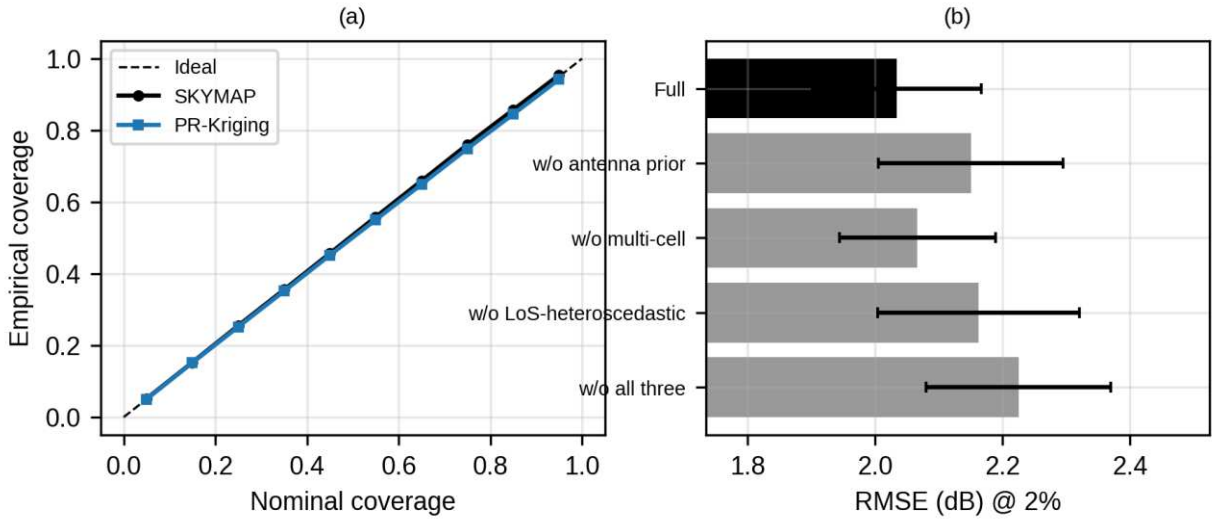

Fig. 5 (a) Reliability diagram for city 1 at 2% sampling, showing empirical versus nominal central-interval coverage (dashed diagonal: ideal). (b) Ablation of SKYMAP components, showing RMSE at 2% sampling (mean ± standard deviation over five cities; the full model is the dark bar)

6.4 Ablation Study

Fig. 5b and Table 7 remove each SKYMAP component at 2% sampling. The LoS-heteroscedastic scale contributes most (+0.13 dB RMSE when removed), followed by the antenna-aware prior (+0.12 dB). Multicell coupling contributes a smaller +0.03 dB. Its gain is limited because every measurement location observes all cells at once, so cross-cell information adds little at the measured points. We expect a larger benefit when cells are measured at different locations, for example with partial neighbour-cell

reporting. Removing all three components gives 2.23 dB, essentially PR-Kriging (2.22 dB), which confirms that the gains come from the proposed components rather than from implementation differences.

6.5 Robustness to Model Mismatch

The SKYMAP kernel mirrors the 3GPP model that generated the data, so a fair evaluation must also test it when that structure is wrong. We regenerated the five cities with no inter-site correlation and a constant shadowing standard deviation of 8 dB in both LoS and NLoS. Table 7 shows that SKYMAP correctly learns $\rho \approx 0.03$, $\kappa \approx 0.05$ and $\sigma_L \approx \sigma_N \approx 8.8$ dB, so its structure collapses gracefully toward a homoscedastic model. However, it underestimates the measurement noise ($\sigma_n = 0.70$ dB versus the true 2 dB), and its RMSE (3.74 dB) is 1.9% worse than that of PR-Kriging (3.67 dB). SKYMAP should therefore be preferred when the 3GPP-style structure is at least approximately valid, which aerial measurement campaigns support [3, 4]. Where the structure is doubtful, fixing σ_n from receiver specifications is a simple safeguard. PINF degrades much more (5.43 dB).

Table 7 Ablation (nominal scenario) and model-mismatch stress test; RMSE (dB) at 2% sampling

Configuration	Nominal scenario	Mismatched scenario
SKYMAP (full)	**2.03 ± 0.13**	3.74 ± 0.15
without antenna-aware prior	2.15 ± 0.15	–
without multicell coupling	2.07 ± 0.12	–
without LoS-heteroscedastic scale	2.16 ± 0.16	–
without all three	2.23 ± 0.15	–
PR-Kriging	2.22 ± 0.12	**3.67 ± 0.15**
PINF	2.91 ± 0.22	5.43 ± 0.11

6.6 Adaptive Sampling: A Negative Result

A natural extension is to use the posterior to choose where the UAV should measure next. We tested three model-driven policies against uniform random sampling, starting from a common random 0.5% set and adding 0.5% per round up to 3%. The policies were space-filling (farthest point), maximum posterior variance, and integrated variance reduction (IVR), the A-optimal criterion of [28]. IVR is computed as $\mathrm{IVR}(c) = \Sigma_b \Sigma_\{\mathbf{x}\in\mathscr{R}\} \Sigma_b(\mathbf{x}, c)^2 / (\sigma_b^2(c) + \sigma_n^2)$ over a reference set $\mathscr{R}$. All selections enforced a minimum spacing of 30 m. Fig. 6 shows that uniform random sampling was best at every budget, reaching 1.71 ± 0.10 dB at 3%, compared with 1.99 ± 0.16 dB for IVR, 2.46 ± 0.18 dB for space-filling and 2.53 ± 0.21 dB for maximum variance. Recalibrating the prior and hyper-parameters only on the initial random subset did not remove the gap in a pilot test on city 1.

We attribute this to a bias–variance effect that is well known in experimental design [29]. Variance-driven criteria trust the model covariance and ignore errors in the mean function, such as prior misfit near LoS/NLoS transitions. Random designs are robust to such misspecification because they sample the whole domain in proportion to its volume, which is exactly how RMSE is measured. We therefore

recommend uniform random sampling within flight constraints for map construction, and we reserve model-driven designs for settings with a well-validated mean function.

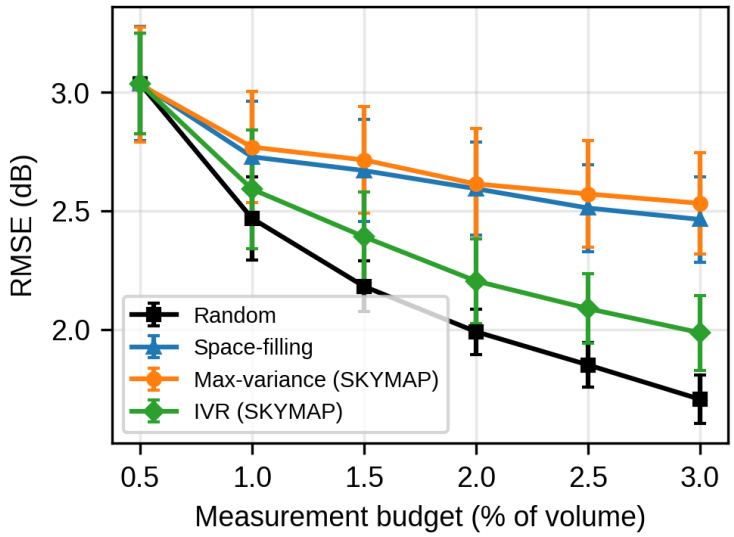

Fig. 6 RMSE versus measurement budget for four sampling policies (mean ± standard deviation over five cities). Uniform random sampling (square markers) outperforms all model-driven policies

6.7 Risk-Weighted Trajectory Planning

We built maps from 2% random samples and planned 40 trajectories per city (200 in total). On average 26.5% of graph nodes are in outage at $\gamma_{th} = 0$ dB, mostly at higher altitudes where LoS interference dominates. Table 8 and Fig. 7 summarise the results. The shortest path spends 37.3% of its length in outage. Planning on a deterministic map helps, giving 16.5% with PR-Kriging and 12.7% with the SKYMAP mean, which reflects the higher node-level outage classification accuracy of SKYMAP (86.1% versus 83.1%). The largest gain comes from using uncertainty: the SKYMAP-risk planner reduces outage to 2.5 ± 0.2%, close to the 1.0% of the oracle, with a path 12% longer than the shortest path. A deterministic planner steers toward cells it believes are just above the threshold, many of which are in fact in outage. The risk-weighted planner instead prices each cell by its probability of outage and keeps a margin wherever the map is uncertain.

Table 8 Trajectory planning with 2% sampling (200 trajectories over five cities)

Planner	Outage along path (%)	Path length (m)	Node outage classification accuracy (%)
Shortest path	37.3 ± 10.2	532 ± 15	–
PR-Kriging (mean map)	16.5 ± 3.4	557 ± 18	83.1
SKYMAP-mean	12.7 ± 3.1	562 ± 15	86.1
SKYMAP-risk (Algorithm 2)	**2.5 ± 0.2**	598 ± 13	86.3
Oracle (true map)	1.0 ± 0.1	555 ± 11	100

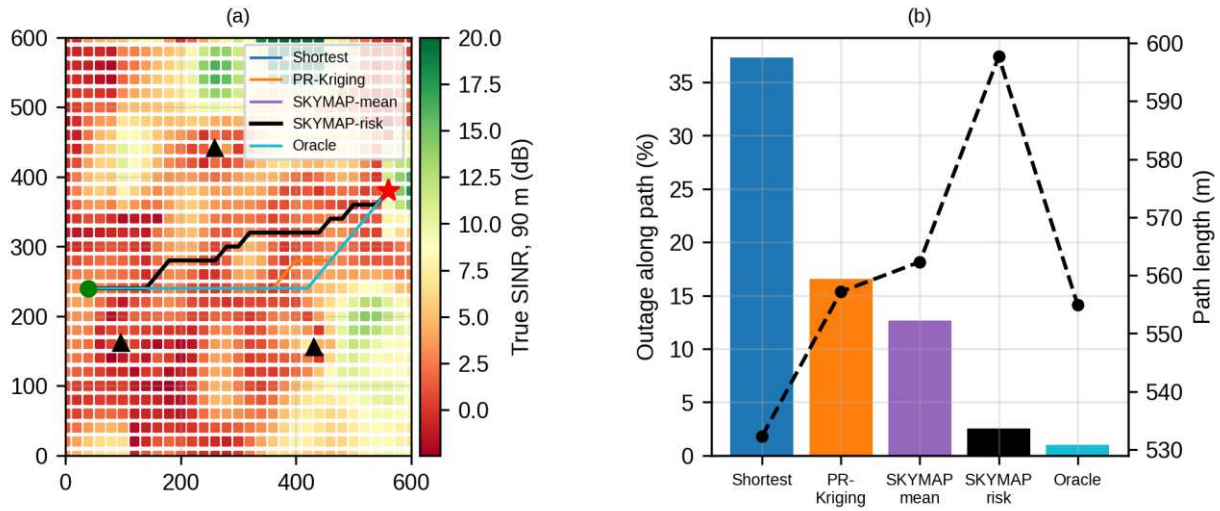

Fig. 7 (a) Example trajectories in city 1 projected onto the 90 m true-SINR layer (circle: start; star: goal; triangles: BSs). (b) Mean outage fraction (bars) and path length (dashed line with markers) per planner over 200 trajectories

6.8 Length-Budgeted Planning with LB-RAP

A fixed risk weight gives no guarantee on path length. With $\lambda = 8$, 55.5% of the 200 queries exceeded a 10% length budget, 17.0% exceeded a 20% budget, and the worst query was 44.6% longer than the shortest path. Table 9 and Fig. 8 evaluate LB-RAP (Algorithm 3) under budgets of 5%, 10% and 20%. LB-RAP satisfied the budget in all 600 query–budget combinations. On SKYMAP maps, it reduced outage to 10.0% with only 3.3% extra length ($\delta = 5\%$) and to 5.4% with 7.6% extra length ($\delta = 10\%$). These are reductions of 73% and 86% relative to the shortest path. With $\delta = 20\%$, LB-RAP matches the fixed-weight planner (2.7% versus 2.5%), so budgets beyond about 13% bring no further benefit in these cities.

The same algorithm applied to the deterministic PR-Kriging map stalls at 16.8 to 19.4% outage whatever the budget. Its hard outage indicator assigns zero cost to cells that are wrongly predicted to be safe, so extra length cannot be spent where it is useful. This result shows that the benefit of LB-RAP depends on calibrated outage probabilities, which is precisely what the SKYMAP posterior supplies. LB-RAP needed 0.06 to 0.11 s per query on one CPU core, which is compatible with in-flight replanning.

Table 9 Length-budgeted planning (Algorithm 3) over 200 queries; extra length is relative to the shortest path

Planner	Budget δ	Outage along path (%)	Extra length (%)	Budget satisfied (%)	Time per query (s)
SKYMAP-risk, fixed $\lambda = 8$	–	2.5 ± 0.2	12.6 ± 2.7	44.5 ($\delta = 10\%$)	0.01
PR-Kriging + LB-RAP	5%	19.4 ± 3.6	2.1 ± 0.5	100	0.05
SKYMAP + LB-RAP	5%	**10.0 ± 3.1**	3.3 ± 0.4	100	0.11
PR-Kriging + LB-RAP	10%	17.2 ± 3.5	3.7 ± 1.1	100	0.03
SKYMAP + LB-RAP	10%	**5.4 ± 1.4**	7.6 ± 0.6	100	0.09
PR-Kriging + LB-RAP	20%	16.8 ± 3.5	4.9 ± 1.7	100	0.02
SKYMAP + LB-RAP	20%	**2.7 ± 0.4**	14.3 ± 1.1	100	0.06

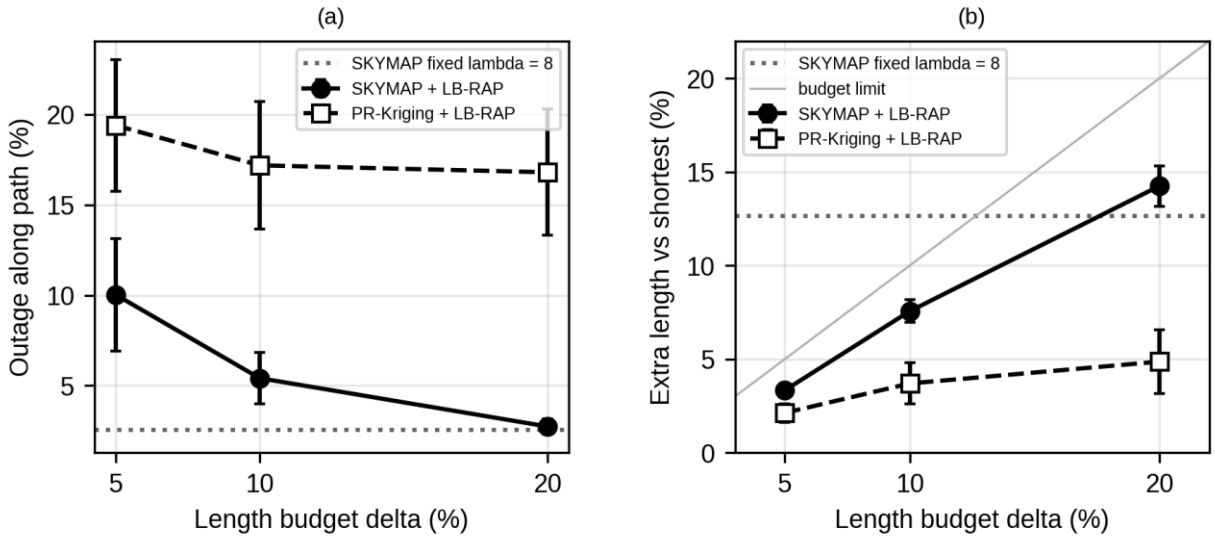

Fig. 8 LB-RAP versus length budget δ: (a) outage along the path and (b) extra length relative to the shortest path (mean ± standard deviation over five cities). Filled circles show SKYMAP + LB-RAP, open squares show PR-Kriging + LB-RAP, the dotted line shows SKYMAP with a fixed $\lambda = 8$, and the thin diagonal in (b) is the budget limit

7 Limitations and Future Work

This study has limitations that should be weighed when interpreting the results.

- **Simulation only.** All results come from simulated cities with 3GPP-based channels. Validation on real UAV measurement data is an essential next step.
- **Scale.** The exact GP scales cubically, so city-scale deployments with many cells will need inducing-point or local approximations.
- **City-model accuracy.** The LoS indicator assumes an accurate 3D city model. Errors in building heights would degrade the prior, and this sensitivity has not yet been quantified.
- **Static map.** The planners treat the map as static and ignore time-varying load, antenna reconfiguration and handover dynamics [5].
- **Adaptive sampling.** Section 6.6 shows that model-driven sampling needs criteria that are robust to errors in the mean function. Developing such criteria remains an open problem.

8 Conclusion

This paper presented SKYMAP, a physics-informed multicell Gaussian process framework that builds 3D radio maps of the low-altitude airspace from sparse measurements and plans UAV trajectories that account for map uncertainty. Placing antenna and LoS physics in the mean and 3GPP-consistent shadowing statistics in the covariance reduced reconstruction error by 10.5 to 13.1% relative to physics-residual kriging and by about half relative to kriging. SKYMAP also outperformed neural radio fields

in the sparse regime, produced well-calibrated uncertainty and recovered interpretable channel parameters. Feeding the posterior outage probability into a shortest-path planner cut the in-flight outage from 37.3% to 2.5%. The proposed LB-RAP planner adds a hard length guarantee and reached 5.4% outage with only 7.6% extra length. The negative findings on model mismatch and adaptive sampling delimit where the approach should and should not be used.

References

1. Zeng, Y., Wu, Q., & Zhang, R. (2019). Accessing from the sky: A tutorial on UAV communications for 5G and beyond. *Proceedings of the IEEE, 107*(12), 2327–2375. https://doi.org/10.1109/JPROC.2019.2952892
2. Mozaffari, M., Saad, W., Bennis, M., Nam, Y.-H., & Debbah, M. (2019). A tutorial on UAVs for wireless networks: Applications, challenges, and open problems. *IEEE Communications Surveys & Tutorials, 21*(3), 2334–2360. https://doi.org/10.1109/COMST.2019.2902862
3. Lin, X., Yajnanarayana, V., Muruganathan, S. D., Gao, S., Asplund, H., Maattanen, H.-L., Bergström, M., Euler, S., & Wang, Y.-P. E. (2018). The sky is not the limit: LTE for unmanned aerial vehicles. *IEEE Communications Magazine, 56*(4), 204–210. https://doi.org/10.1109/MCOM.2018.1700643
4. 3GPP. (2017). *Enhanced LTE support for aerial vehicles* (Technical Report TR 36.777, Version 15.0.0). 3rd Generation Partnership Project.
5. Arul Dalton, G., Louis, A. B. V., Ramachandran, A., & Savija, J. (2025). Modified deep learning model in proactive decision-making for handover management in 5G. *Concurrency and Computation: Practice and Experience, 37*(9–11), e70023. https://doi.org/10.1002/cpe.70023
6. Bi, S., Lyu, J., Ding, Z., & Zhang, R. (2019). Engineering radio maps for wireless resource management. *IEEE Wireless Communications, 26*(2), 133–141.
7. Zeng, Y., Chen, J., Xu, J., Wu, D., Xu, X., Jin, S., Gao, X., Gesbert, D., Cui, S., & Zhang, R. (2024). A tutorial on environment-aware communications via channel knowledge map for 6G. *IEEE Communications Surveys & Tutorials, 26*(3), 1478–1519.
8. Romero, D., & Kim, S.-J. (2022). Radio map estimation: A data-driven approach to spectrum cartography. *IEEE Signal Processing Magazine, 39*(6), 53–72.
9. Levie, R., Yapar, Ç., Kutyniok, G., & Caire, G. (2021). RadioUNet: Fast radio map estimation with convolutional neural networks. *IEEE Transactions on Wireless Communications, 20*(6), 4001–4015. https://doi.org/10.1109/TWC.2021.3054977
10. Chen, J., Yatnalli, U., & Gesbert, D. (2017). Learning radio maps for UAV-aided wireless networks: A segmented regression approach. In *Proceedings of the IEEE International Conference on Communications (ICC)*. IEEE.
11. Zhang, S., & Zhang, R. (2021). Radio map-based 3D path planning for cellular-connected UAV. *IEEE Transactions on Wireless Communications, 20*(3), 1975–1989.
12. Zeng, Y., Xu, X., Jin, S., & Zhang, R. (2021). Simultaneous navigation and radio mapping for cellular-connected UAV with deep reinforcement learning. *IEEE Transactions on Wireless Communications, 20*(7), 4205–4220.
13. Archana, T., Aravind, T., Malini, A. H., & Kalaivani, C. T. (2026). Energy-aware adaptive obstacle avoidance based on meta-reinforcement learning with segmentation for UAV trajectory planning. *International Journal of Aeronautical and Space Sciences, 27*(1), 598–618. https://doi.org/10.1007/s42405-025-00950-6

14. Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2020). NeRF: Representing scenes as neural radiance fields for view synthesis. In *Computer Vision – ECCV 2020*. Springer.
15. Zhao, X., An, Z., Pan, Q., & Yang, L. (2023). NeRF²: Neural radio-frequency radiance fields. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM.
16. Cressie, N. (1993). *Statistics for spatial data* (Rev. ed.). Wiley.
17. Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press.
18. Manju, S., Prabha, M., Farithkhan, A., Sivagurunathan, G., & Michael Mahesh, K. (2025). Decentralized control design for UAV swarms communication. *Discover Applied Sciences, 7*(2), 131. https://doi.org/10.1007/s42452-024-06408-w
19. Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., & Ng, R. (2020). Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems 33*.
20. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics, 378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045
21. Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems 30*.
22. Al-Hourani, A., Kandeepan, S., & Lardner, S. (2014). Optimal LAP altitude for maximum coverage. *IEEE Wireless Communications Letters, 3*(6), 569–572. https://doi.org/10.1109/LWC.2014.2342736
23. 3GPP. (2022). *Study on channel model for frequencies from 0.5 to 100 GHz* (Technical Report TR 38.901, Version 17.0.0). 3rd Generation Partnership Project.
24. Gudmundson, M. (1991). Correlation model for shadow fading in mobile radio systems. *Electronics Letters, 27*(23), 2145–2146. https://doi.org/10.1049/el:19911328
25. Álvarez, M. A., Rosasco, L., & Lawrence, N. D. (2012). Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning, 4*(3), 195–266.
26. Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik, 1*, 269–271. https://doi.org/10.1007/BF01386390
27. ITU-R. (2012). *Propagation data and prediction methods required for the design of terrestrial broadband radio access systems operating in a frequency range from 3 to 60 GHz* (Recommendation P.1410-5). International Telecommunication Union.
28. Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research, 9*, 235–284.
29. Box, G. E. P., & Draper, N. R. (1959). A basis for the selection of a response surface design. *Journal of the American Statistical Association, 54*(287), 622–654. https://doi.org/10.1080/01621459.1959.10501525

Statements and Declarations

Funding The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Competing Interests The author has no relevant financial or non-financial interests to disclose.

Author Contributions G. Arul Dalton conceived and designed the study, developed the methodology and software, performed the experiments and analysis, and wrote and approved the final manuscript.

Data Availability No measurement data were used. All datasets analysed in this study were generated synthetically by the code in the repository https://github.com/GArulDalton/skymap-uav-radio-maps, which also contains the raw result files behind every table and figure.

Code Availability The complete source code is openly available at https://github.com/GArulDalton/skymap-uav-radio-maps , under the MIT licence.

Ethics Approval and Consent Not applicable. This study did not involve human participants or animals.