drones

MDPI

Article

Environmental Perception via Propagation-Equivalent Reconstruction from Opportunistic Cellular Signals

Zhiang Bian [1], Hu Lu [1,2,*], Cheng Zhang [1], Zhisen Wang [1], Hongcheng Li [1], Xin He [1], Liangdong Wang [1] and Fan Hu [1]

[1] Information and Navigation School, Air Force Engineering University, Xi'an 710077, China; bianzzaa@163.com (Z.B.); zhangch21@foxmail.com (C.Z.); 2019031031@cauc.edu.cn (Z.W.); lihongcheng22@alumni.nudt.edu.cn (H.L.); hexin_lyn@163.com (X.H.); 15998709030@163.com (L.W.); hufan285412@outlook.com (F.H.)
[2] National Key Laboratory of Unmanned Aerial Vehicle Technology, Xi'an 710077, China
* Correspondence: sdkmsdn@163.com

Highlights

What are the main findings?

- Protocol fields in cellular downlinks carry an intrinsic power reference. This anchors the loss scale and resolves the ambiguity between transmit power and environmental attenuation, turning scalar RSRP from an underdetermined mixture into a constrained environmental observable.
- Coarse site bearings and multi-altitude diversity separate horizontal blockage from vertical clearance, making height-dependent spatial reconstruction possible without prior building geometry.

What are the implication of the main findings?

- Signals of opportunity, long confined to receiver positioning and timing, extend naturally to environmental perception: the same downlink observations that constrain navigation also carry recoverable spatial structure.
- Ambient cellular infrastructure provides an environmental sensing resource for UAVs, removing the need for dedicated illuminators, network-side cooperation, or prior building maps.

Abstract

Low-altitude unmanned aerial vehicles (UAVs) need environmental evidence for perception and navigation, yet dedicated onboard sensing and prior maps may be unavailable. Existing cellular infrastructure offers persistent signals of opportunity, but a passive UAV observes only scalar RSRP, which conflates transmit power, sector response, and environmental loss. We formulate UAV environmental perception as passive, protocol-assisted sensing using non-cooperative commercial cellular downlinks and propose a protocol-anchored reconstruction of a propagation-equivalent virtual radio environment map (vREM). Decodable nominal reference power and path-loss normalization set the power scale, coarse site bearings constrain antenna directions, and multi-altitude UAV trajectories excite height dependence. A hard height-class Beer–Lambert model recovers occupancy support and a height proxy without building geometry at inference. In a controlled multi-altitude simulation based on Xi'an Bell Tower geometry, the method achieved a held-out $RMSE_{RSRP}$ of 5.24 dB, compared with 15.70 dB for NoProtocol; all metrics improved across 200 paired repetitions. A separate ground-based field collection yielded 1518 observations from 82 LTE/NR sources at 41 sites and reduced relative-power mapping RMSE from 7.27 to 3.41 dB. The receiver uses public broadcast fields locally, while the commercial transmitters remain outside the sensing and inversion loop.

Check for updates

Academic Editor: Neng Ye

Received: 31 July 2026
Revised: 3 September 2026
Accepted: 9 September 2026
Published: 15 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Keywords: UAV environmental perception; cellular signals of opportunity; non-cooperative sensing; virtual radio environment maps (vREMs); radio tomographic imaging (RTI); reference signal received power (RSRP)

1. Introduction

Low-altitude UAV perception is commonly built around dedicated onboard sensors, while cellular links are treated mainly as communication or navigation resources. In cities, these functions share a physical cause: buildings shape traversable space and imprint obstruction on cellular downlink propagation [1–3]. Prior building maps and network-side engineering data may also be unavailable to a third-party platform. Extracting propagation structure from signals already present in the airspace would therefore provide a complementary environmental evidence layer without adding a dedicated illuminator.

The key requirement is non-cooperative observability. A third-party receiver observes scalar RSRP, in which true transmit power, path loss, sector response, and building blockage share one link budget. Network-side ISAC, known-power RTI, and operator-maintained Radio Maps or CKMs can use controlled transmitters or complete engineering parameters [4–6]. Map-assisted prediction uses building geometry as an input. The present setting therefore calls for a formulation that converts passive downlink power into interpretable environmental evidence from decodable protocol fields and multi-altitude geometry, while the transmitter remains outside the sensing loop.

Throughout this paper, we therefore describe the architecture as passive, protocol-assisted sensing using non-cooperative commercial cellular downlinks. "Passive" dictates that the receiver emits no probing signal; "protocol-assisted" dictates that it decodes public synchronization and broadcast fields; and "non-cooperative" dictates that the commercial transmitter remains outside the sensing and inversion loop. The receiver uses the standardized synchronization, reference signal, and broadcast information chain to establish a reproducible observation scale, while the transmitter continues its ordinary service.

Signals-of-opportunity navigation and Radio SLAM show that non-cooperative radio signals can support receiver or source state estimation [7,8]. These navigation methods target position, clock, or velocity, while propagation enters mainly as multipath structure or NLOS error. Radio Maps instead predict location-indexed received power or coverage fields through interpolation and completion [9,10], and CKMs organize path loss, shadowing, and other channel descriptors as spatial priors for communication decisions [6,11]. Both extend the spatial representation of radio states. The unresolved step for UAV perception is to treat the propagation environment itself, rather than the receiver state or coverage intensity, as the reconstruction target.

Virtual obstacle radio mapping moves from field prediction to environmental inversion. It infers propagation-defined obstruction from RSS or RSRP attenuation, making the environment the sensing target [12–14]. These inversions typically fix the loss scale through known transmit power, antenna patterns, or controlled link geometry. The third-party cellular setting requires recovering propagation structure from decodable broadcast power, deterministic path loss, coarse site bearings, and multi-altitude trajectories, while explaining independently held-out RSRP.

We construct three observation anchors from passively available information. Decodable nominal reference power and deterministic path loss fix the power scale. Coarse site bearings restrict unknown sector antennas to a low-dimensional directional subspace. Multi-altitude UAV trajectories provide geometric excitation that separates horizontal obstruction from vertical clearance. With these anchors, non-cooperative RSRP becomes a

height-gated Beer–Lambert inverse problem with low-dimensional transmitter-side nuisance parameters. Its output is a vREM containing occupancy support and a height proxy. We define a propagation-equivalent vREM as a latent occupancy and height-proxy field whose primary constraint is to explain independently held-out RSRP, dominant blockage, and vertical-clearance effects under the stated link model. The equivalence is defined by these propagation observables, so the map supplies a physically interpretable propagation and connectivity evidence layer for candidate route screening and connectivity-aware planning.

Contributions

1. We formulate propagation environment perception for a passive UAV receiver as a non-cooperative cellular inverse problem. Nominal reference power, path-loss normalization, coarse site bearings, and multi-altitude trajectories jointly constrain transmitter–environment ambiguity.
2. We develop a propagation-equivalent vREM estimator matched to these observation conditions. A continuous height-gated forward model, hard height-class inversion, source-specific biases, and constrained antenna updates recover occupancy support and a height proxy.
3. We provide complementary evidence through controlled multi-altitude UAV simulation and ground-based real-road collection. Two-hundred paired repetitions show stable anchoring gains, while 1518 field observations support protocol recovery and multi-source evidence formation in an operational receiver chain.

The remainder of this paper is organized as follows. Section 2 reviews related work, and Section 3 formulates the system model. Section 4 analyzes observability, while Section 5 presents the estimator. Section 6 reports the experiments, and Section 7 concludes the paper.

Notation: Italic lowercase letters, bold lowercase letters, bold uppercase letters, and calligraphic letters denote scalars, vectors, matrices, and sets, respectively. The set of real numbers is $\mathbb{R}$. The transpose, Moore–Penrose pseudoinverse, and rank of $\mathbf{A}$ are $\mathbf{A}^\top$, $\mathbf{A}^\dagger$, and $\mathrm{rank}(\mathbf{A})$. Its largest and smallest singular values are $\sigma_{\max}(\mathbf{A})$ and $\sigma_{\min}(\mathbf{A})$. We use $\|\mathbf{x}\|_2$ for the Euclidean norm and $|\mathcal{G}|$ for set cardinality. The expectation and the indicator function are $\mathbb{E}[\cdot]$ and $\mathbf{1}[\cdot]$. The Hadamard product is $\odot$, and $\mathrm{proj}_{\mathcal{C}}(\cdot)$ is Euclidean projection onto $\mathcal{C}$.

2. Related Work

2.1. Passive Observation of Non-Cooperative Cellular Downlinks

Integrated sensing and communication (ISAC) typically adopts a network-side perspective. It designs base-station waveforms, beams, resources, or cooperative links for sensing [4,15,16]. This paradigm suits operator-led low-altitude surveillance and joint communication–sensing deployments. The observation link is configurable, synchronization and resources are schedulable, and transmitter states and antenna patterns are the system parameters. A passive third-party UAV receiver instead monitors existing downlink reference signals, with waveform, resource, and transmitter-state control remaining on the network side. It therefore relies solely on publicly available observations. The present study uses this receiver-side information to formulate the non-cooperative inverse problem.

Passive radar instead uses broadcast and communication signals as illuminators of opportunity for target detection and tracking. Recent studies have used 5G NR signals for range–Doppler detection of vehicles and UAVs [17,18]. This paradigm shares our non-cooperative use of cellular downlinks. Its targets, however, are the positions and velocities of moving scatterers rather than static obstruction and propagation-loss fields. Passive radar also relies on target-reflected echoes. We use direct-path RSRP attenuation

caused by penetration or blockage, which is a different propagation component of the same downlink signal.

Cellular signals-of-opportunity navigation and Radio SLAM use non-dedicated signals to constrain receiver trajectories, clocks, or source positions [7,8]. In that navigation setting, environment-related structure acts as an anchor for estimating the platform state. The present setting reverses the input–output relation: UAV poses are treated as observed inputs, while propagation-equivalent occupancy and height are reconstructed outputs. Environmental perception therefore requires direct separation of distance loss, sector angle, building blockage, and altitude dependence in RSRP.

Adjacent passive radar and signal-of-opportunity studies similarly decode standardized protocol structure without transmitter coordination [7,19–21].

2.2. Radio Environment Reconstruction from RSS to Spatial Fields

Radio tomographic imaging (RTI) provides a direct line-integral model from RSS to spatial loss fields. Wilson and Patwari used Beer–Lambert absorption to recover a two-dimensional loss field from multiple RSS links [5]. Subsequent work addressed noise models [22,23], fade-level priors [24], and Bayesian regularization [25]. These models attribute RSS variation approximately to spatial absorption. In passive third-party RSRP, the residual also contains uncalibrated power, sector sidelobes and nulls, source-specific offsets, and altitude-dependent height gating. Separating these terms is therefore central to the present formulation: direct two-dimensional RTI would otherwise combine antenna, power, and blockage effects in one loss field.

Radio Map interpolation and deep completion instead approach coverage estimation as a prediction task. RadioUNet and deep completion autoencoders can produce accurate power maps from training data or environmental inputs [9,10]. Their strength is function approximation and held-out power prediction. Their latent representations need not separate distance loss, transmit power, antenna angle, or blockage. We retain these methods as predictive baselines for held-out RSRP. A propagation-equivalent vREM additionally constrains prediction to interpretable environmental variables. In practice, 3GPP TS 37.320 specifies Stage-2 minimization-of-drive-test (MDT) radio measurements [26]. MDT-based coverage prediction uses building and geographic features to interpolate RSRP [27]. Its objective is network optimization, whereas the present study targets third-party non-cooperative sensing.

A channel knowledge map (CKM) organizes location-tagged channel descriptors into a reusable regional database [6,11]. Closely related CKM construction methods require dense position labels, historical channel samples, network maintenance, ray tracing, or environmental models [28,29]. This paradigm supports regional channel database construction and maintenance. Our inputs are single-platform passive RSRP, nominal protocol power, coarse site priors, and UAV poses. The output is an equivalent environmental variable constrained by a propagation model. Representative RSS-to-field methods usually assume known, controlled, or network-maintained transmitter states. Our separation problem instead includes both unknown transmit power and unknown antenna response.

Open geographic data and ray tracing have also advanced Radio Map construction. Geo2SigMap converts OpenStreetMap building footprints into Sionna ray-tracing scenes and achieved 6.04 dB RSRP prediction error on measured LTE data [30]. The first path-loss Radio Map prediction challenge introduced a community benchmark with building geometry as input [31]. These methods assume known building geometry and material properties. They predict coverage given environmental geometry, whereas we recover propagation-equivalent structures when that geometry is unavailable. The distinct inputs and outputs make the two problem settings complementary.

2.3. Virtual Obstacles and vREMs

The closest line of work is the virtual obstacle radio mapping conducted by Chen and their collaborators. Air-to-ground channel reconstruction was used to infer propagation-defined obstacle fields from RSS [12]. Liu and Chen extended this approach to joint least-squares estimation of three-dimensional multi-class virtual obstacles and propagation parameters [13]. W. Chen and J. Chen then introduced multiple knife-edge diffraction and geometry-assisted deep learning to represent diffraction and scattering [14]. These studies show that RSS contains inversion information about obstruction, material class, or environmental semantics beyond coverage intensity.

Virtual obstacle models define propagation-based environmental variables, but they generally use stronger observation conditions. These include known absolute transmit power, controlled geometry, environmental supervision, or known and simplified antenna patterns. Non-cooperative cellular RSRP combines environmental loss, nominal-to-true power mismatch, sector patterns, source-specific biases, and measurement noise in one scalar. Existing virtual obstacle methods take power anchors and antenna models as inputs. We instead include nominal protocol power, low-dimensional antennas, and source-specific biases in one estimation problem. Nominal protocol power replaces an engineering absolute-power anchor, while a constrained antenna block replaces a known pattern. The resulting vREM is propagation-equivalent because its primary role is to explain held-out RSRP under the constrained forward model.

Recent work has also moved toward sparse 3D REM recovery and propagation-structure-aware learning. Wang et al. used sparse Bayesian learning and optimized sampling to recover 3D spectrum maps under channel shadowing [32]; Chen and Chen used diffraction and obstruction structure to learn a virtual environment without an explicit city map [33]. These studies reinforce that radio measurements can contain environmental information. Our formulation combines passive cellular protocol anchors, incomplete sector parameters, source-specific biases, and multi-altitude UAV excitation in one third-party receiver model.

Recent surveys make a complementary distinction between source-aware forward prediction and source-agnostic inverse reconstruction, and identify uncalibrated receivers, non-unique solutions, and multipath as open issues in the latter [34]. Recent systems illustrate the two sides of this distinction: FERMI uses a hybrid propagation model with cooperative robot data collection [35]; CKM-enabled dynamic REM reconstruction assumes that transmitter locations, powers, and radiation patterns are available for lightweight forward generation [36]; and RadioFormer3D estimates low-altitude 3D Radio Maps from sparse measurements with weak supervision [37]. These studies provide context for the input–output distinction of the present work. Our receiver passively reads public cellular broadcast fields, estimates transmitter-side nuisance terms, and reconstructs a propagation-equivalent obstruction field rather than a source-aware coverage map.

3. System Model and Problem Formulation

We formulate non-cooperative downlink vREM reconstruction as a structured tomographic inverse problem with low-dimensional nuisance parameters. Four anchors define the model: nominal protocol power fixes the power scale; a source-specific bias absorbs nominal-to-actual link budget mismatch; a constrained sector antenna block restricts angular uncertainty; and a height-gated Beer–Lambert model converts multi-altitude observations into height constraints. A single attenuation class avoids bilinear scale ambiguity, and in the dB domain, antenna gain and bias enter as additive nuisance terms. The resulting observation equation matches the implemented height-class estimator.

Figure 1 summarizes the observation-to-inversion chain. Multi-altitude passive RSRP provides geometric excitation, while nominal protocol power and path-loss normalization form a calibrated residual. Low-dimensional antenna and bias blocks absorb transmitter-side nuisance variation. The height-gated tomography block outputs occupancy support and a height proxy. This section presents the full model, and Section 5 gives the estimator.

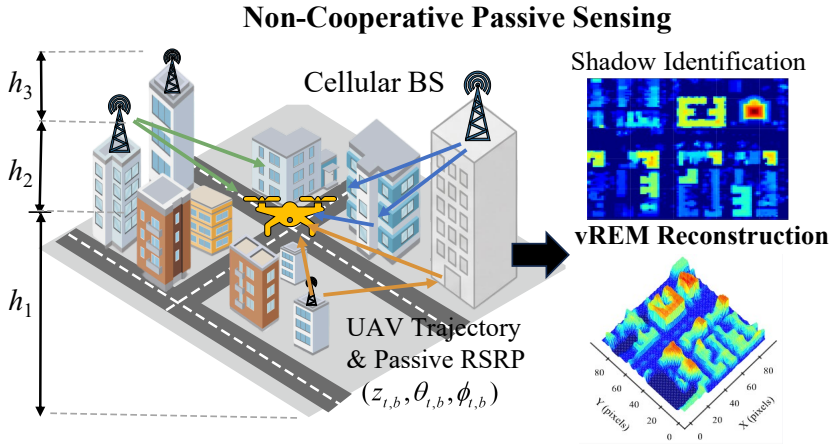


Figure 1. Protocol-anchored vREM reconstruction from non-cooperative cellular downlinks. The colors on the right schematically represent height.

3.1. Link Model and Calibrated Residual

A UAV equipped with a passive signals-of-opportunity receiver traverses an urban region, $\mathcal{D} \subset \mathbb{R}^2$, at $h_{\min} \le h \le h_{\max}$. Decodable LTE/NR downlink sources form the finite set $\mathcal{B} = \{1, \ldots, B\}$. These non-cooperative sources may occupy rooftops, facades, or pole-mounted sites, and their engineering parameters may be incomplete. The UAV samples the channel at epochs $\mathcal{T} = \{1, \ldots, T\}$. The subset observed at epoch t is $\mathcal{B}_t \subseteq \mathcal{B}$. We partition the horizontal sensing region into N square cells of side length Δ_g. Cell $g \in \mathcal{G} = \{1, \ldots, N\}$ has center $\mathbf{x}^{(g)} \in \mathbb{R}^2$ and support $\Omega_g \subset \mathcal{D}$. The UAV and source positions are $\mathbf{x}_t = [x_t, y_t, h_t]^\top$ and $\mathbf{s}_b = [x_b, y_b, h_b]^\top$. The total number of observation tuples is

$$M = \sum_{t \in \mathcal{T}} |\mathcal{B}_t| \le TB. \tag{1}$$

Each non-cooperative downlink ray, $\ell_{t,b}$, intersects horizontal cells. Cell occupancy, discrete height class, and a fixed attenuation scale explain the calibrated environmental loss. The same residual also contains the effective sector gain $A_b(\theta_{t,b}, \phi_{t,b}; \boldsymbol{\psi}_b)$. Protocol power, site bearing, and multi-altitude geometry must therefore separate spatial blockage from low-dimensional antenna direction.

We define the operational sampling envelope as $\mathcal{E} = \mathcal{D} \times [h_{\min}, h_{\max}]$ and require $\mathbf{x}_t \in \mathcal{E}$. The observation coordinate, $\mathbf{x}_t$, is referenced to the receiving antenna phase center. In the controlled experiment, $\mathcal{D}$ is a 1000 m $\times$ 1000 m horizontal inversion window discretized at 20 m, and inference is not extrapolated outside this window. At this macro-scale resolution, the UAV is represented as a point receiver: vehicle body geometry and safety clearance are not components of the vREM state. They remain separate flight system constraints to be applied by the trajectory planner and near-field sensing together with the propagation evidence supplied by the vREM.

For each observation pair, (t, b), the slant range, horizontal range, azimuth, and elevation from $\mathbf{s}_b$ to $\mathbf{x}_t$ are

$$d_{t,b} = \|\mathbf{x}_t - \mathbf{s}_b\|_2, \tag{2}$$

$$\rho_{t,b} = \sqrt{(x_t - x_b)^2 + (y_t - y_b)^2}, \tag{3}$$

$$\theta_{t,b} = \operatorname{atan2}(y_t - y_b,\, x_t - x_b), \tag{4}$$

$$\phi_{t,b} = \arctan\big((h_t - h_b)/\rho_{t,b}\big). \tag{5}$$

The propagation path is $\boldsymbol{\ell}_{t,b}(\lambda) = \mathbf{s}_b + \lambda(\mathbf{x}_t - \mathbf{s}_b)$ for $\lambda \in [0,1]$, with arc-length element $d_{t,b}\, d\lambda$. For cell g with horizontal support Ω_g, the intersection $\Lambda_{t,b,g} = \{\lambda \in [0,1] : \Pi_{xy}\boldsymbol{\ell}_{t,b}(\lambda) \in \Omega_g\}$ defines the chord traversed by the ray. The horizontal ray–cell intersection length is

$$a_{t,b,g} = d_{t,b}\, \cos\phi_{t,b}\, \big|\Lambda_{t,b,g}\big| \tag{6}$$

The ray height within cell g is

$$h^{\text{ray}}_{t,b,g} = h_b + \bar{\lambda}_{t,b,g}(h_t - h_b), \tag{7}$$

where $\bar{\lambda}_{t,b,g} = \frac{1}{2}(\inf \Lambda_{t,b,g} + \sup \Lambda_{t,b,g})$. Stacking $\{a_{t,b,g}\}$ by row gives the sparse tomography matrix $\mathbf{A} \in \mathbb{R}^{M\times N}_{\geq 0}$. It has $\mathcal{O}(M\sqrt{N})$ nonzero entries and satisfies $\sum_g a_{t,b,g} = d_{t,b}\cos\phi_{t,b}$.

3GPP specifies received power measurements separately from downlink power configuration. LTE RSRP and NR SS-RSRP are the receiver-side linear averages calculated over specified reference signal resource elements [38,39]. Our nominal power anchor instead comes from passively decodable broadcast RRC and system information: for LTE, SystemInformationBlockType2 provides `referenceSignalPower`, the configured CRS EPRE [40]; for NR, SIB1 provides `ss-PBCH-BlockPower`, the configured SSS EPRE, with within-block EPRE relations specified by TS 38.213 [41,42]. These operations require only broadcast decoding and do not place the transmitter inside the sensing loop.

$$P_b^{\text{ref}} \equiv \begin{cases} P^{\text{cfg}}_{b,\text{CRS-EPRE}}, & \text{LTE}, \\ P^{\text{cfg}}_{b,\text{SSS-EPRE}}, & \text{NR}. \end{cases} \tag{8}$$

The strong anchor set contains only sources with complete decodable fields. Sources with missing fields can be treated as weak anchors under a wider δ_b prior. The quantitative experiments use strong anchors by default. Each anchor provides configured CRS EPRE (LTE) or configured SSS EPRE (NR). A source-specific bias absorbs remaining differences between this broadcast configuration and the actual link budget, including transmitter implementation, hardware calibration, and receiver chain gain. We denote this configured nominal reference power by P_b^{ref} in dBm. The per-resource-element received power then follows the link budget

$$z_{t,b} = P_b^{\text{ref}} + \delta_b + A_b(\theta_{t,b}, \phi_{t,b}; \boldsymbol{\psi}_b) - PL_0 - L_{\text{env}} + w_{t,b}, \tag{9}$$

Here, $\delta_b \in [\delta_{\min}, \delta_{\max}]$ is the *source-specific RSRP bias*. It absorbs calibration drift, slowly varying temperature-dependent gain, and small unmodeled differences between nominal and actual reference signal power. A finite interval breaks the additive source–environment gauge freedom and prevents unbounded drift. The effective sector gain, $A_b(\cdot)$, is parameterized by $\boldsymbol{\psi}_b = [\alpha_b, \beta_b]^\top$. Angular correction, $\boldsymbol{\psi}_b$, enters through the nonlinear gain, $A_b(\cdot; \boldsymbol{\psi}_b)$, while scalar bias, δ_b, is an additive source-specific offset; Section 4 shows

that these terms have different observability roles. The deterministic distance–frequency baseline path loss is

$$PL_0(d,f) \;=\; \zeta_0 + 10n_0\log_{10}(d/d_0) + 20\log_{10}(f/f_0), \tag{10}$$

which follows the close-in parameterization of 3GPP TR 38.901 [43,44]. The term L_{env} is excess environmental loss. The term $w_{t,b}$ aggregates shadowing, fast fading, and measurement noise. We define the protocol-calibrated residual

$$y_{t,b} \;\triangleq\; P_b^{\text{ref}} - z_{t,b} - PL_0(d_{t,b}, f_b), \tag{11}$$

Substitution into (9) gives the observation equation

$$y_{t,b} \;=\; L_{\text{env}} - A_b(\theta_{t,b}, \phi_{t,b}; \boldsymbol{\psi}_b) - \delta_b + e_{t,b}, \tag{12}$$

where $e_{t,b} = -w_{t,b}$. On unobstructed links, the residual is therefore dominated by low-dimensional antenna gain and source-specific bias. Both A_b and δ_b must be constrained in the same estimation problem.

The residual $e_{t,b}$ contains three components. The log-normal shadowing term $e_{t,b}^{\text{shad}} \sim \mathcal{N}(0,\sigma_s^2)$ follows the Gudmundson exponential covariance kernel [45]

$$\text{Cov}[e_{t,b}^{\text{shad}}, e_{t',b'}^{\text{shad}}] = \sigma_s^2\,\mathbf{1}[b=b']\,e^{-\|\mathbf{x}_t-\mathbf{x}_{t'}\|_2/d_c}, \tag{13}$$

The fast-fading residual is $e_{t,b}^{\text{fade}} \sim \mathcal{N}(0,\sigma_f^2)$, where $\sigma_f^2 \ll \sigma_s^2$ because RSRP averaging suppresses fast fading. The heavy-tailed measurement component is

$$e_{t,b}^{\text{meas}} \sim (1-\epsilon_o)\mathcal{N}(0,\sigma_m^2) + \epsilon_o\mathcal{N}(0,\sigma_o^2), \tag{14}$$

The controlled simulation explicitly includes measurement outliers. Paired Monte Carlo repetitions test the squared-loss estimator under this perturbation. Stacking $y_{t,b}$ into $\mathbf{y} \in \mathbb{R}^M$ and using the source indicator matrix $\mathbf{D} \in \{0,1\}^{M\times B}$ gives

$$\mathbf{y} \;=\; \mathbf{L}_{\text{env}}(\mathcal{M}) - \mathbf{a}(\boldsymbol{\psi}) - \mathbf{D}\boldsymbol{\delta} + \mathbf{e}, \tag{15}$$

where $\boldsymbol{\delta} = [\delta_1,\ldots,\delta_B]^\top \in \mathbb{R}^B$ and $\boldsymbol{\psi} = [\boldsymbol{\psi}_1^\top,\ldots,\boldsymbol{\psi}_B^\top]^\top \in \mathbb{R}^{2B}$.

The Gudmundson kernel couples shadowing samples from the same source within correlation distance d_c. Section 4 uses the ideal Mahalanobis metric $\boldsymbol{\Sigma}_e^{-1}$ to analyze information projection. The implemented estimator instead uses squared residuals, nonnegative sparsity, and bounded source-specific biases. Two-hundred paired repetitions assess empirical stability under correlated shadowing, measurement noise, and outliers.

3.2. Antenna and Bias Model

A fully flexible source-specific angular pattern is structurally unobservable along one UAV trajectory. It could reinterpret many spatial blockage signatures as pattern fluctuations. We therefore constrain A_b to the 3GPP TR 36.873/38.901 sector template [2,43,46,47]. Each source has only two angular parameters, $\boldsymbol{\psi}_b = [\alpha_b, \beta_b]^\top$. We define $\Delta\theta_{t,b} = \text{wrap}_\pi(\theta_{t,b} - \alpha_b)$ and positive downward as $\Delta\phi_{t,b} = \phi_{t,b} + \beta_b$. Data generation and candidate evaluation use the hard-clipped template

$$\begin{aligned} u_H(\Delta\theta) &= 12(\Delta\theta/\Theta_{3\mathrm{dB}})^2, \\ u_V(\Delta\phi) &= 12(\Delta\phi/\Phi_{3\mathrm{dB}})^2, \\ D_{HV}^{\mathrm{hard}} &= \min\{D_m, u_H + u_V\}, \\ A_b^{\mathrm{hard}} &= G_{\max} - D_{HV}^{\mathrm{hard}}. \end{aligned} \tag{16}$$

where $G_{\max}$ is peak gain and $D_m = 30\,\mathrm{dB}$ is the attenuation cap. Equation (9) shows the synthetic RSRP, and each candidate in (35) uses $A_b = A_b^{\mathrm{hard}}$. The local Fisher analysis in Section 4 replaces only the hard clipping with the differentiable proxy

$$D_{HV}^{\mathrm{sm}} = \mathrm{softmin}_{D_m}(u_H + u_V), \qquad A_b^{\mathrm{sm}} = G_{\max} - D_{HV}^{\mathrm{sm}}, \tag{17}$$

where $\mathrm{softmin}_{D_m}(u) = -\tau \log(e^{-u/\tau} + e^{-D_m/\tau})$ and $\tau = 1.0$. As $\tau \to 0$, A_b^{sm} converges to the hard template in (16). With $\omega = \mathrm{sigmoid}((D_m - u_H - u_V)/\tau)$, the proxy has gradients

$$\frac{\partial A_b^{\mathrm{sm}}}{\partial \alpha_b} = \omega \frac{24\Delta\theta_{t,b}}{\Theta_{3\mathrm{dB}}^2}, \tag{18}$$

$$\frac{\partial A_b^{\mathrm{sm}}}{\partial \beta_b} = -\omega \frac{24\Delta\phi_{t,b}}{\Phi_{3\mathrm{dB}}^2}. \tag{19}$$

The weight ω approaches 1 in the unclipped quadratic region and approaches 0 when the attenuation cap is active. The smooth form supports only local derivative analysis. It does not generate synthetic data or enter the discrete grid solver. The feasible set of $\boldsymbol{\psi}_b$ is the closed convex box, $\mathcal{A}_b = [\alpha_b^{\mathrm{srv}} - \Delta_\alpha, \alpha_b^{\mathrm{srv}} + \Delta_\alpha] \times [\beta_{\min}, \beta_{\max}]$. The bias lies in the independent box, $\mathcal{D}_b = [\delta_{\min}, \delta_{\max}]$. Here, α_b^{srv} is a coarse sector bearing from public data, third-party surveying, or a pre-mission site survey. The tolerance Δ_α covers installation error, with $\Delta_\alpha \approx 5^\circ$ representing a practical value.

A two-parameter sector template (α_b, β_b) is used: half-power beamwidths $(\Theta_{3\mathrm{dB}}, \Phi_{3\mathrm{dB}})$ follow the standardized operator deployments, and additional pattern degrees of freedom would absorb spatial blockage directions and worsen the information separation analyzed in Section 4. Survey accuracy near 5° confines α_b to a neighborhood of $\alpha_b^\star$, supporting the local analysis in Section 4.

3.3. Height-Gated vREM

We represent the urban environment as a 2.5D absorbing medium. Cell-wise occupancy, $o : \mathcal{D} \to [0,1]$, describes propagation-blockage support. Virtual height, $h : \mathcal{D} \to [0, H_{\max}]$, provides a height-gating proxy. We fix a single attenuation scale, $\eta_g \equiv \eta_0$. Free cell-wise attenuation would multiply occupancy and create scale ambiguity. Fixing η_0 focuses recovery on occupancy support and height information supported by non-cooperative observations. The resulting vREM explains independently held-out RSRP. Its propagation-defined support and height proxy complement material and cadastral layers by supplying radio-derived environmental evidence.

In the dB domain, the continuous height-gated Beer–Lambert blockage loss along ray $\boldsymbol{\ell}_{t,b}$ is

$$L_{\mathrm{blk}}^{\mathrm{cont}} = \int_0^1 o\,\eta\,q\big(h - h_{t,b}^{\mathrm{ray}}(\lambda)\big)\,d_{t,b}\cos\phi_{t,b}\,d\lambda, \tag{20}$$

where $q(u) = \mathrm{sigmoid}(\kappa_h u)$. The controlled simulation applies this continuous gate to true occupancy and height when generating environmental loss. The same model defines multi-altitude excitation and the information directions in Section 4. Inversion does not access true heights. It instead uses the hard-gate relaxation over the discrete candidate heights below.

Discretizing (o, h, η) as a piecewise constant over $\{\Omega_g\}$ and applying (6) and (7) gives

$$L_{\text{blk}} = \sum_{g \in \mathcal{G}} a_{t,b,g}\, o_g\, \eta_g\, q_{t,b,g}, \tag{21}$$

The discretized continuous gate is

$$q_{t,b,g} = \text{sigmoid}\big(\kappa_h (h_g - h^{\text{ray}}_{t,b,g})\big). \tag{22}$$

The matrix form is

$$\mathbf{L}_{\text{blk}} = \big(\mathbf{A} \odot \mathbf{Q}^{(\mathbf{h})}\big)\,(\mathbf{o} \odot \boldsymbol{\eta}). \tag{23}$$

The implemented single-class closure sets $\boldsymbol{\eta} = \eta_0 \mathbf{1}$, reducing (23) to $\eta_0(\mathbf{A} \odot \mathbf{Q}^{(\mathbf{h})})\mathbf{o}$. As $\kappa_h \to \infty$, $q_{t,b,g} \to \mathbf{1}[h_g > h^{\text{ray}}_{t,b,g}]$, which gives a *hard-height gate*. In this binary limit, rays far below h_g undergo near-complete attenuation, while rays far above h_g pass freely. The resulting contrast supports coarse height classification in Section 6. In differentiable regions of the continuous model, the Jacobian is

$$\partial L_{\text{blk}} / \partial o_g = \eta_g\, a_{t,b,g}\, q_{t,b,g}, \tag{24}$$

$$\partial L_{\text{blk}} / \partial h_g = \eta_g\, a_{t,b,g}\, o_g\, \kappa_h\, q_{t,b,g}(1 - q_{t,b,g}). \tag{25}$$

The factor $q(1-q) \in (0, \frac{1}{4}]$ in (25) is largest when a ray grazes the top of a virtual obstacle. Multi-altitude measurements constrain cell–height, h_g, most strongly in this geometry.

$$L_{\text{env}} = L_{\text{blk}}. \tag{26}$$

3.4. Structured Inversion Objective

Let $\mathcal{H}_{\text{cls}} = \{h_k\}_{k=1}^{K}$ denote the discrete height classes, and index observations by $i = (t, b)$. In the $\kappa_h \to \infty$ limit, restricting h_g to $\mathcal{H}_{\text{cls}}$ gives the sparse hard-gated class matrix $\mathbf{H} \in \mathbb{R}_{\geq 0}^{M \times NK}$

$$H_{i,(g,k)} = \eta_0 a_{i,g}\, \mathbf{1}[h_k > h^{\text{ray}}_{i,g}]. \tag{27}$$

If each occupied cell selects one height class, let $s_{g,k}$ denote its occupancy intensity in that class. Then,

$$L_i^{\text{hard}} = \sum_{g,k} H_{i,(g,k)} s_{g,k}, \quad s_{g,k} \geq 0, \quad \|\mathbf{s}_g\|_0 \leq 1. \tag{28}$$

Here, $\mathbf{s}_g = [s_{g,1}, \ldots, s_{g,K}]^\top$. Column (g,k) represents the loss contribution of cell g at height class h_k across all links. Direct cell-wise class selection in (28) is combinatorial. We relax it to $\mathbf{x} \geq 0$ and control activation using nonnegative ℓ_1 regularization, yielding (29).

The inverse uses the hard-gate discrete height-class dictionary in (27); the continuous sigmoid gate in (21) and (22) is used for forward generation and supplies differentiable directions for the information analysis in Section 4.

Given the antenna parameters, $\boldsymbol{\psi}$, and source-specific biases, $\boldsymbol{\delta}^{(j)}$, the environmental block solves

$$\begin{aligned} \mathbf{x}^{(j+1)} &= \arg\min_{\mathbf{x} \geq 0} \tfrac{1}{2}\|\mathbf{H}\mathbf{x} - \mathbf{y}_\psi^{(j)}\|_2^2 \\ &\quad + \lambda_1 \|\mathbf{x}\|_1, \\ \mathbf{y}_\psi^{(j)} &= \mathbf{y} + \mathbf{a}(\boldsymbol{\psi}) + \mathbf{D}\boldsymbol{\delta}^{(j)}. \end{aligned} \tag{29}$$

where $x_{g,k}$ is the nonnegative intensity for a cell–height pair. The clipped source-specific bias update is

$$\delta_b^{(j+1)} = \mathrm{proj}_{[-8,8]}\left[\frac{1}{|\mathcal{I}_b|}\sum_{i\in\mathcal{I}_b}\left((\mathbf{H}\mathbf{x}^{(j+1)})_i - y_i - a_i(\boldsymbol{\psi}_b)\right)\right], \tag{30}$$

The environmental and bias blocks alternate. After the fixed-scale Gaussian smoothing of each class score, we denote the result by $\widetilde{x}_{g,k}$. The final output is

$$\begin{aligned}\hat{o}_g &= \min\left(1, \sum_k \widetilde{x}_{g,k}\right),\\ \hat{h}_g &= h_{\arg\max_k \widetilde{x}_{g,k}}, \qquad \hat{o}_g > 0.02.\end{aligned} \tag{31}$$

The height proxy is undefined for all remaining cells. Thus, the estimator variables, outputs, and experimental metrics correspond directly. The threshold $\tau^\star$ applies to $\hat{o}_g$ only during controlled structural evaluation and does not enter fitting. The estimator uses the nonnegative ℓ_1 objective in (29); no ℓ_2 penalty is used.

4. Observability Analysis

This section explains why the three observation anchors help separate the environment from transmitter-side nuisance parameters. For the continuous analytical proxy, local Fisher information under Gaussian shadowing characterizes the information projection caused by marginalizing antenna parameters and biases; held-out prediction, geometric coverage diagnostics, and paired repetitions evaluate the implemented discrete estimator. We also introduce coverage conditions which are checkable directly from trajectory and site geometry.

Under the single-class attenuation convention in Section 3.3, the unknown parameter vector is $\boldsymbol{\theta} = [\mathbf{o}^\top, \mathbf{h}^\top, \boldsymbol{\psi}^\top, \boldsymbol{\delta}^\top]^\top \in \mathbb{R}^{2N+3B}$. The local analysis evaluates derivatives of the smooth proxy in (17). Its observation mean is $\boldsymbol{\mu}^{\mathrm{sm}}(\boldsymbol{\theta}) = \mathbf{L}_{\mathrm{env}}(\mathcal{M}) - \mathbf{a}^{\mathrm{sm}}(\boldsymbol{\psi}) - \mathbf{D}\boldsymbol{\delta}$. The Jacobian $\mathbf{J}(\boldsymbol{\theta}) = \partial\boldsymbol{\mu}^{\mathrm{sm}}/\partial\boldsymbol{\theta}^\top \in \mathbb{R}^{M\times(2N+3B)}$ has the block decomposition $\mathbf{J} = [\mathbf{J}_\mathbf{o}, \mathbf{J}_\mathbf{h}, \mathbf{J}_{\boldsymbol{\psi}}, \mathbf{J}_{\boldsymbol{\delta}}]$. Equations (18), (24) and (25) give these blocks. Under the Gaussian-shadowing limit, $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_e)$, the Fisher information matrix (FIM) is

$$\mathcal{I}(\boldsymbol{\theta}) \;=\; \mathbf{J}(\boldsymbol{\theta})^\top \boldsymbol{\Sigma}_e^{-1} \mathbf{J}(\boldsymbol{\theta}). \tag{32}$$

After marginalizing the angular antenna parameters, $\boldsymbol{\psi}$, and source-specific biases, $\boldsymbol{\delta}$, the Schur complement gives information about $\mathbf{o}$:

$$\mathcal{I}_{\mathbf{o}|\boldsymbol{\psi},\boldsymbol{\delta}} = \mathbf{J}_\mathbf{o}^\top \boldsymbol{\Sigma}_e^{-1/2} (\mathbf{I} - \mathbf{P}_\mathrm{a}) \boldsymbol{\Sigma}_e^{-1/2} \mathbf{J}_\mathbf{o}, \tag{33}$$

where $\mathbf{J}_\mathrm{a} = [\mathbf{J}_{\boldsymbol{\psi}}, \mathbf{J}_{\boldsymbol{\delta}}] \in \mathbb{R}^{M\times 3B}$ collects the angular and bias columns of the Jacobian, and

$$\mathbf{P}_\mathrm{a} = (\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathrm{a})(\mathbf{J}_\mathrm{a}^\top \boldsymbol{\Sigma}_e^{-1} \mathbf{J}_\mathrm{a})^\dagger (\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathrm{a})^\top$$

is the orthogonal projection onto $\mathrm{col}(\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathrm{a})$. The Moore–Penrose pseudoinverse accommodates antenna or bias directions that are weakly excited by a particular trajectory. If the nuisance block has full column rank, (33) reduces to the standard Schur complement with an ordinary inverse.

Lemma 1 (Bound on antenna-induced information loss)**.** *The information-loss matrix* $\mathbf{L}_{\mathrm{loss}} = \mathbf{J}_\mathbf{o}^\top \boldsymbol{\Sigma}_e^{-1/2} \mathbf{P}_\mathrm{a} \boldsymbol{\Sigma}_e^{-1/2} \mathbf{J}_\mathbf{o}$ *has a rank of* $3B$ *at most. Angular parameters contribute* $2B$ *dimensions*

and biases contribute B dimensions. Its range lies within the column space of $\mathbf{J}_\mathbf{o}^\top \boldsymbol{\Sigma}_e^{-1}\mathbf{J}_\mathrm{a}$*. Thus, marginalizing antennas and biases removes information about* $\mathbf{o}$ *only within a subspace of a dimension of* $3B$ *at most in* $\mathbb{R}^N$*. For any direction,* $\mathbf{v}\in\mathbb{R}^N$*, the loss magnitude satisfies*

$$\mathbf{v}^\top \mathbf{L}_\mathrm{loss}\mathbf{v} \;\le\; \|\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathbf{o}\mathbf{v}\|_2^2 \cos^2 \angle(\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathbf{o}\mathbf{v}, \boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathrm{a}), \tag{34}$$

where $\angle(\cdot,\cdot)$ *is the angle between the first vector and the column space of the second argument.*

Proof. The rank of $\mathbf{P}_\mathrm{a}$ is $\mathrm{rank}(\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathrm{a}) \le 3B$, so $\mathrm{rank}(\mathbf{L}_\mathrm{loss}) \le 3B$. For any $\mathbf{v}$, $\mathbf{v}^\top\mathbf{L}_\mathrm{loss}\mathbf{v} = \|\mathbf{P}_\mathrm{a}\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathbf{o}\mathbf{v}\|_2^2$ is the squared projection of $\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathbf{o}\mathbf{v}$ onto $\mathrm{col}(\boldsymbol{\Sigma}_e^{-1/2}\mathbf{J}_\mathrm{a})$. This identity gives (34). □

Equation (34) replaces the heuristic "$\mathcal{O}(B/M)$ information fraction" with a principal-angle bound between occupancy- and antenna-sensitive subspaces. The bound is informative when this angle remains separated from zero. The geometric coverage conditions below address this requirement.

This rank bound characterizes the directional information loss caused by antenna and bias marginalization. A small principal angle, $\angle(\mathbf{J}_\mathbf{o}, \mathbf{J}_\mathrm{a})$, increases the condition number, $\kappa(\mathcal{I}_{\mathbf{o}|\boldsymbol{\psi},\delta})$, of the effective information matrix and weakens the angular–spatial separation. The site-bearing constraints in Section 3.2 and diagnostics (C1)–(C3) limit this confounding.

Geometric Coverage Diagnostics

Protocol anchoring yields useful environmental information when observations sufficiently excite antenna directions, cross-source paths, and height gates. We therefore use three coverage diagnostics computed directly from UAV and base-station geometry. They are scene-dependent design targets that guide pre-flight geometry checks and provide a direct link between trajectory design and reconstruction quality:

(C1) Per-source azimuth span:The design target is trajectory-induced azimuth coverage of approximately $2\Theta_{3\mathrm{dB}}$ for the surveyed bearing, α_b^{srv}, for each source, b. The empirical azimuth histogram and the number of retained observations are reported together.

(C2) Multi-source occupancy coverage: For each cell, g, in high-occupancy support, $\mathcal{S}$, we record whether at least two physical-source rays traverse the cell and whether their maximum angular separation reaches the scene's $\Theta_{3\mathrm{dB}}$ scale. One pass through the sparse tomography matrix, $\mathbf{A}$, gives the coverage and angular-pass rates.

(C3) Multi-altitude diversity: We record whether the trajectory contains two or more altitude layers and whether their total span is comparable to the median propagation-equivalent obstruction-height scale. UAV altitude records verify this diagnostic directly.

C1 converts coarse site bearings into sector-angle excitation. C2 checks cross-source path constraints on each spatial cell, while C3 tests whether the trajectory spans typical obstruction heights. Together, these diagnostics turn protocol, site, and trajectory anchors into a checkable observation design. The matched scans in Section 6.6 quantify how each design factor changes prediction, support, and height metrics, so the same diagnostics can be used to configure a new route. Appendix A reports their measured coverage in the Xi'an scene.

5. Solution Algorithm

5.1. Three-Block Alternating Estimation

The implemented estimator alternates among environmental, source-specific bias, and low-dimensional antenna blocks. The environmental block applies nonnegative ℓ_1 tomography over discrete height classes. It jointly produces nonzero occupancy support

and a height proxy under a fixed attenuation scale, η_0. The bias block absorbs source-level differences between nominal protocol power and the actual link budget. The antenna block performs constrained discrete updates within a preset domain, defined by coarse site bearings and nominal downtilts. Height-class initialization forms environmental structure before updating low-dimensional transmitter-side nuisance parameters. This order prevents occupancy and height from suppressing each other at zero initialization. The value $\eta_0 = 1.5\,\mathrm{dB/m}$ is the fixed single-class scale used for the formal experiment; its role and sensitivity are reported in Appendix C. The 20 m grid is the selected resolution–coverage–computation compromise for the 1000 m scene.

Implemented estimation chain:

(1) Set $\boldsymbol{\psi}^{(0)} = \boldsymbol{\psi}^{\mathrm{srv}}$ and $\delta^{(0)} = \mathbf{0}$. Construct protocol-calibrated target $\mathbf{y} + \mathbf{a}(\boldsymbol{\psi}^{(0)})$ from the surveyed antenna state.
(2) With antennas fixed, construct the hard height-class matrix using (27). Solve (29) for environmental intensity, $\mathbf{x}$, using nonnegative accelerated proximal gradients.
(3) Update each source bias using (30). Alternate environmental and bias blocks three times to obtain the initial height-class field.
(4) For each source, b, update the antenna state within the survey-constrained candidate domain, $\mathcal{G}_b$, by

$$\boldsymbol{\psi}_b^{(r+1)} = \arg\min_{\boldsymbol{\psi}_b \in \mathcal{G}_b} \frac{1}{|\mathcal{I}_b|} \times \sum_{i \in \mathcal{I}_b} \left(y_i - (\mathbf{H}\mathbf{x}^{(r)})_i + a_i(\boldsymbol{\psi}_b) + \delta_b^{(r)} \right)^2 \tag{35}$$

(5) Re-solve the environmental and bias blocks using updated $\boldsymbol{\psi}^{(r+1)}$. Perform steps (4) and (5) for three outer environmental–antenna updates.
(6) Obtain $(\hat{\mathbf{o}}, \hat{\mathbf{h}})$ from (31). Generate held-out RSRP predictions using the link budget in (9).

These six steps cover height-class matrix construction, nonnegative ℓ_1 environmental updates, clipped bias updates, constrained antenna updates, outer resolving, and output collapse. This sequence matches the code path used for the main table, spatial maps, and 200 Monte Carlo repetitions.

5.2. Implementation Conventions

Height classes follow target-city building scales at the 20 m grid; the single attenuation class fixes the occupancy scale, and site survey accuracy defines the feasible antenna and bias domains. These settings remain fixed across all experiments.

6. Experiments and Results

The experiments test three claims. First, that protocol and geometric anchoring improve held-out RSRP prediction under non-cooperative reception. Second, that the gains persist across random splits, biases, noise, outliers, and controlled model mismatch. Third, that a real ground-based passive receiver chain can form mappable multi-source environmental evidence. Held-out $\mathrm{RMSE}_{\mathrm{RSRP}}$ is the primary metric. The IoU*, F1*, and RMSE_h metrics diagnose a propagation-equivalent structure in controlled scenes.

Single-run and field RMSE are reported to two decimals, Monte Carlo means and standard deviations to one decimal, heights and distances to one decimal, and bounded proportions to two decimals.

6.1. Scene and Observation Protocol

The controlled experiment covers a 1000 m × 1000 m area around Xi'an Bell Tower. Building footprints and heights come from the public CMAB dataset [48], and roads come from the OSM [49]. After common projection, the scene contains 356 height-annotated buildings and 93 roads, providing reproducible urban geometry for non-cooperative RSRP inversion. The cellular deployment follows 3GPP three-dimensional channel and low-altitude UAV coverage models [1,2,43,46]. It contains 12 sites and 24 decodable sectors. The simulator uses complete transmitter parameters to generate RSRP. The estimator receives only passive RSRP, nominal protocol reference power, coarse site and antenna priors, and UAV poses.

Each synthetic RSRP sample follows (9): protocol power and source bias set the transmit scale, (10) provides distance–frequency loss, the 3GPP template gives the true sector gain, and the continuous sigmoid gate with true occupancy and height determines environmental loss, followed by Gaussian perturbations and outliers. During inversion, Prop solves (29) using the hard height-class dictionary in (27) without accessing true occupancy or height.

A 20 m grid gives $N = 2500$ horizontal cells. The estimation set samples 900 horizontal positions across seven altitude layers at 5/15/30/45/60/80/100 m. The holdout set samples 220 disjoint horizontal positions across three different layers at 20/50/85 m. They retain $M_{\text{est}} = 123{,}395$ and $M_{\text{test}} = 13{,}422$ passive observation tuples, respectively. The results are reported for the tested 5–100 m flight envelope. The discrete height set is $\mathcal{H}_{\text{cls}} = \{5, 8, 11, 14, 18, 22, 28, 36, 45, 55, 65\}$ m. Appendix A reports diagnostics for sector azimuth span, cross-source coverage, and height excitation. Experimental conclusions rely on held-out prediction, structural diagnostics, and protocol-front-end ablation.

Figure 2 summarizes the public urban geometry, controlled cellular sector deployment, and road-constrained UAV sampling support.

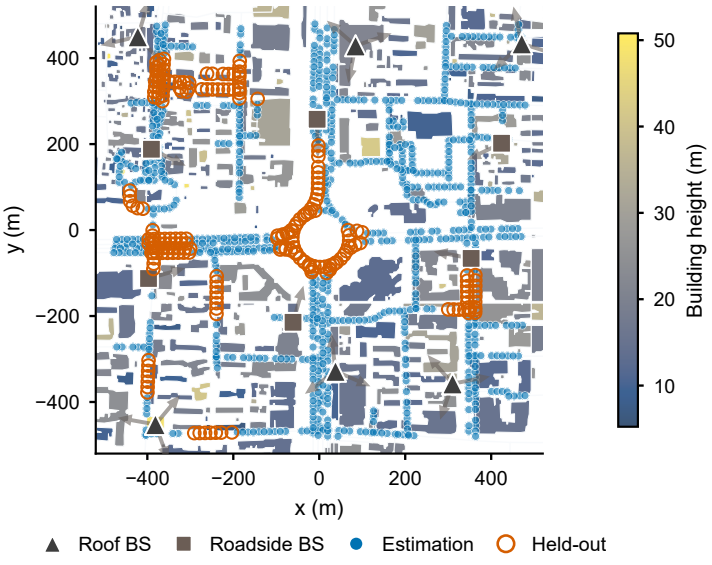


Figure 2. Xi'an Bell Tower scene and sampling. Arrows indicate sector azimuths.

Figure 3 shows how protocol and survey calibration affect distance trends. The calibrated loss target, $\tilde{y}_{t,b} = y_{t,b} + A_b^{\text{surv}}$, adds surveyed antenna correction to the residual in (11). Across $M_{\text{est}} = 123{,}395$ tuples, pooled absolute Pearson correlation with link distance drops from 0.19 (raw RSRP) to 0.09 (calibrated target), and per-source correlation decreases for 22 of 24 sources. Calibration reduces aggregate distance dependence while

preserving spatial loss variation for structured inversion; remaining surveyed antenna error and source biases are jointly explained by the subsequent model.

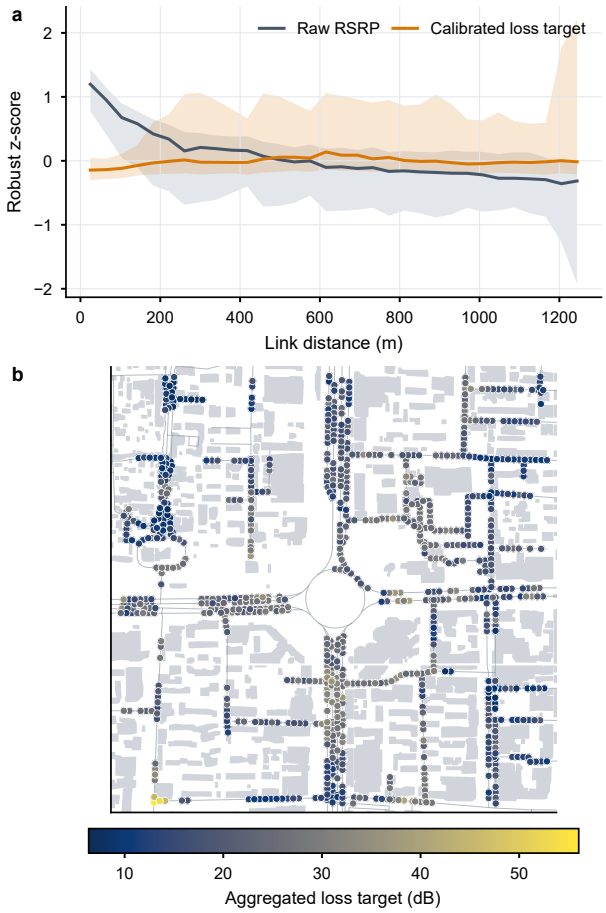


Figure 3. RSRP loss targets before and after protocol calibration. (**a**) Distance dependence; (**b**) spatial distribution of the aggregated calibrated loss target.

6.2. Baselines and Metrics

The comparison includes four method groups, none of which receive true environment or true antenna inputs. The first group contains the Radio Map prediction baselines. These are per-source inverse-distance weighting (IDW), a raw-RSRP Kriging or Gaussian-process regression (GPR) predictor, and a coordinate multilayer perceptron (MLP) without protocol reference power. They predict only held-out RSRP and therefore do not report vREM structure metrics. Kriging–Raw denotes the local Kriging/GPR baseline in the single-run main table. RBF-Gauss denotes per-source Gaussian-kernel or RBF smoothing of raw RSRP in the Monte Carlo analysis. The second group contains classical two-dimensional tomography diagnostics, R-Tikh and SART [5,50]. Both use attenuation scale $\eta_0 = 1.5\,\mathrm{dB/m}$ and output a two-dimensional loss or support proxy. The third group contains virtual obstacle (VO) and environmental-semantics (EnvSem)-compatible baselines. These baselines adapt the ideas in [13,14] to our non-cooperative observation setting. VO-Compat and EnvSem-Compat therefore use an observable link budget front end built from nominal protocol power and path-loss normalization. The fourth group contains the combined anchoring ablation and Prop. NoProtocol constructs a loss target from raw RSRP. It omits protocol power, deterministic path loss, surveyed antenna correction, and constrained low-dimensional antenna updates. Prop uses the protocol-calibrated residual, hard height-class vREM, source-specific biases, and constrained low-dimensional antenna updates.

Let i denote a retained passive observation tuple, (t, b). VO-Compat uses the protocol front end to construct

$$\tilde{y}_i = P_{b_i}^{\mathrm{ref}} + A_{b_i}^{\mathrm{surv}} - PL_{0,i} - z_i. \tag{36}$$

Here, $A_{b_i}^{\mathrm{surv}}$ comes only from the surveyed site pattern, and $P_{b_i}^{\mathrm{ref}}$ is the passively decodable nominal protocol power. VO-Compat solves

$$\min_{\mathbf{x}\geq 0,\delta} \frac{1}{2}\|\mathbf{H}\mathbf{x} - \tilde{\mathbf{y}} - \mathbf{S}\delta\|_2^2 + \lambda\|\mathbf{x}\|_1, \tag{37}$$

where $\mathbf{H}$ is the height-gated ray matrix built from UAV and source geometry and candidate height classes. Vector $\mathbf{x}$ contains cell–height virtual obstacle intensities, $\mathbf{S}$ is the source indicator matrix, and δ_b is the residual source-specific bias. Its prediction is

$$\hat{z}_i^{\mathrm{PVO}} = P_{b_i}^{\mathrm{ref}} + \hat{\delta}_{b_i} + A_{b_i}^{\mathrm{surv}} - PL_{0,i} - (\mathbf{H}\hat{\mathbf{x}})_i. \tag{38}$$

EnvSem-Compat then fits per-source residual Kriging to $z_i - \hat{z}_i^{\mathrm{PVO}}$:

$$\hat{r}_* = \mathbf{k}_*^\top(\mathbf{K} + \nu\mathbf{I})^{-1}\mathbf{r}, \qquad \mathbf{K}_{mn} = \exp\left(-\frac{\|\mathbf{p}_m - \mathbf{p}_n\|_2^2}{2\ell^2}\right), \tag{39}$$

The final prediction is $\hat{z}_* = \hat{z}_*^{\mathrm{PVO}} + \hat{r}_*$. This comparison places VO and EnvSem ideas under the same observable link budget conditions. It measures the contribution of structured vREM with joint constrained antenna and bias estimation in the same Xi'an case. All baseline configurations were fixed before formal holdout evaluation and fitted on the complete estimation set.

VO-Compat and EnvSem-Compat are compatible controls implemented with our non-cooperatively observable inputs. Appendix B summarizes the non-default estimator and baseline settings.

All quantitative baselines use consistent observable inputs. Training inputs exclude constraints from undetected links. Ray matrices use only UAV and source geometry, without true environmental loss or height gates. Results involving true antenna information only assess sensitivity to antenna priors.

The primary metric is $\mathrm{RMSE}_{\mathrm{RSRP}}$ on disjoint held-out trajectories:

$$\mathrm{RMSE}_{\mathrm{RSRP}} = \sqrt{\frac{1}{M_{\mathrm{test}}}\sum_{(t,b)\in\mathrm{test}}\left(\hat{z}_{t,b} - z_{t,b}\right)^2}, \tag{40}$$

Here, $\hat{z}_{t,b}$ is the model-predicted RSRP from recovered single-class closure $(\hat{\mathbf{o}}, \hat{\mathbf{h}}, \eta_0, \hat{\boldsymbol{\psi}}, \hat{\boldsymbol{\delta}})$. This metric directly evaluates held-out UAV-link prediction through the recovered propagation-equivalent map.

The *secondary* metrics are diagnostic only. Let $\mathcal{S}^\star = \{g : o_g^\star > 0.5\}$ be the true support and $\widehat{\mathcal{S}}_\tau = \{g : \hat{o}_g > \tau\}$ be the thresholded recovered support. The implementation uses an evenly spaced threshold set, $\mathcal{T}_\tau$, and defines

$$\tau^\star = \arg\max_{\tau\in\mathcal{T}_\tau} \frac{|\widehat{\mathcal{S}}_\tau \cap \mathcal{S}^\star|}{|\widehat{\mathcal{S}}_\tau \cup \mathcal{S}^\star|}, \tag{41}$$

$$\mathrm{IoU}^\star = \frac{|\widehat{\mathcal{S}}_{\tau^\star} \cap \mathcal{S}^\star|}{|\widehat{\mathcal{S}}_{\tau^\star} \cup \mathcal{S}^\star|}, \tag{42}$$

$$\mathrm{F1}^\star = \frac{2|\widehat{\mathcal{S}}_{\tau^\star} \cap \mathcal{S}^\star|}{|\widehat{\mathcal{S}}_{\tau^\star}| + |\mathcal{S}^\star|}. \tag{43}$$

The starred scores provide an oracle upper envelope for structural diagnosis. Deployment-oriented evaluation additionally uses a fixed midpoint rule ($\tau = 0.50$) and a ground-truth-free Otsu rule computed from the predicted occupancy-score histogram [51]:

$$\tau_{\mathrm{Otsu}} = \arg\max_{\tau} \ \omega_0(\tau)\omega_1(\tau)\big[\mu_0(\tau) - \mu_1(\tau)\big]^2. \tag{44}$$

For any frozen threshold, τ, let $\mathcal{D}_\tau = \{g : o_g^\star > 0.5, \hat{o}_g \geq \tau, \hat{h}_g < \infty\}$ denote the detected reference-occupied cells with finite height estimates. We retain the original height error only as a conditional diagnostic,

$$\mathrm{RMSE}_h^{\mathrm{cond}}(\tau) = \sqrt{\frac{1}{|\mathcal{D}_\tau|} \sum_{g \in \mathcal{D}_\tau} (\hat{h}_g - h_g^\star)^2}. \tag{45}$$

This metric evaluates a propagation-equivalent height-gating proxy. Its denominator is always reported because it excludes missed cells.

To remove successful-cell selection from the height audit, define the effective recovered height

$$\tilde{h}_g(\tau) = \begin{cases} \hat{h}_g, & \hat{o}_g \geq \tau \text{ and } \hat{h}_g < \infty, \\ 0, & \text{otherwise}, \end{cases} \tag{46}$$

and set $h_g^\star = 0$ on reference-empty cells. We then report

$$\mathrm{RMSE}_h^{\mathrm{occ}}(\tau) = \sqrt{\frac{1}{|\mathcal{S}^\star|} \sum_{g \in \mathcal{S}^\star} \big(\tilde{h}_g(\tau) - h_g^\star\big)^2}, \tag{47}$$

$$\mathrm{RMSE}_h^{\mathrm{grid}}(\tau) = \sqrt{\frac{1}{N} \sum_{g=1}^{N} \big(\tilde{h}_g(\tau) - h_g^\star\big)^2}. \tag{48}$$

The first metric includes every missed reference-occupied cell; the second also penalizes recovered height on false-positive empty cells. Precision, recall, and IoU are reported alongside both height metrics.

6.3. Controlled-Scene Height Correspondence and Full-Support Audit

The height correspondence audit uses the public CMAB `pred_h_r` attribute rasterized at the same 20 m grid, with 824 occupied and 1676 empty cells. Reference heights generate controlled observations but are not provided to the estimator. The primary audit freezes an Otsu threshold from the predicted occupancy-score histogram before accessing reference support.

Table 1 summarizes the height audit. Under the ground-truth-free Otsu rule, Prop detects 636 of 824 reference-occupied cells. Recovered and reference heights have $r = 0.83$ on these detected cells, with a conditional RMSE of 4.7 m, an MAE of 3.4 m, and a bias of +2.0 m. When all occupied cells are retained and misses are assigned zero height, the descriptive correlation is $r = 0.50$, the bias is -2.7 m, and $\mathrm{RMSE}_h^{\mathrm{occ}} = 10.2$ m. Including false positives gives $\mathrm{RMSE}_h^{\mathrm{grid}} = 6.6$ m. NoProtocol yields $r_{\mathrm{det}} = 0.40$, $\mathrm{RMSE}_h^{\mathrm{occ}} = 14.9$ m, and $\mathrm{RMSE}_h^{\mathrm{grid}} = 11.6$ m. Thus, Prop reduces full-grid height map error by 43.1% under the same ground-truth-free threshold rule. These results establish correspondence between the recovered propagation-equivalent height field and the controlled scene reference while evaluating every occupied cell in the global penalty statistics.

Table 1. Controlled scene height correspondence and full-support audit.

Method	Rule	TP	FP	FN	Precision	Recall	IoU	r_{det}	$\mathrm{RMSE}_h^{\mathrm{cond}}$	$\mathrm{RMSE}_h^{\mathrm{occ}}$	$\mathrm{RMSE}_h^{\mathrm{grid}}$
NoProtocol	Otsu	586	136	238	0.81	0.71	0.61	0.40	12.5	14.9	11.6
Prop	Otsu	636	50	188	0.93	0.77	0.73	0.83	4.7	10.2	6.6

All height errors are in meters. r_{det} is the descriptive Pearson correlation across detected reference-occupied cells; grid cells are spatially related. Otsu reads only predicted scores.

6.4. Quantitative Comparison

The central result in Table 2 is the mechanistic gain from the observation front end. Prop achieved a held-out $\mathrm{RMSE}_{\mathrm{RSRP}}$ of 5.24 dB, compared with 15.70 dB for NoProtocol (a 66.6% reduction), 13.68 dB for Kriging–Raw, 12.71 dB for MLP-RawCoord, and 5.81/5.70 dB for VO-Compat/EnvSem-Compat. This matched-model, matched-data ablation shows that protocol and geometric anchoring create a more stable environmental inversion target from raw RSRP. The IoU*, F1*, and conditional RMSE_h metrics diagnose agreement between the recovered latent field and controlled reference structure; the non-oracle full-support audit is reported in Table 1 and Figure 4.

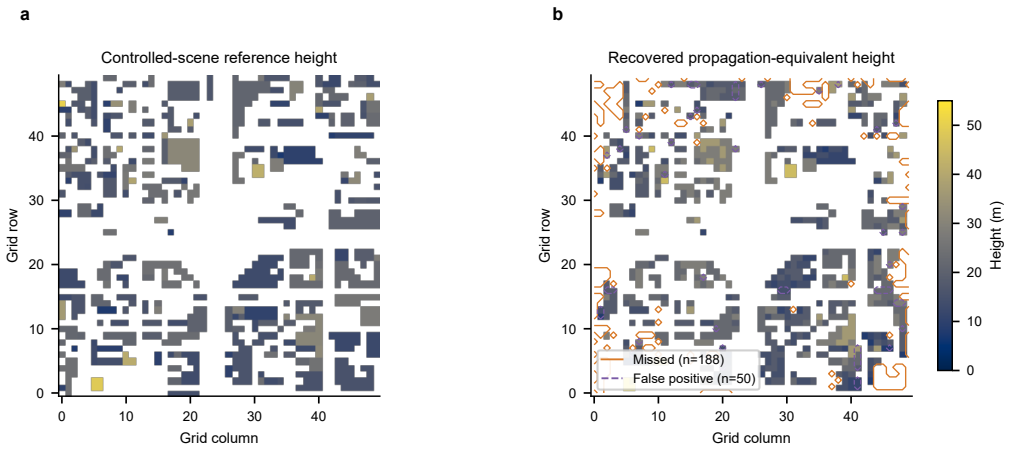


Figure 4. Controlled scene height correspondence and support audit. (**a**) CMAB reference height; (**b**) Prop height proxy with missed and false-positive support contours.

Table 2. Controlled simulation around Xi'an Bell Tower.

Method	$\mathrm{RMSE}_{\mathrm{RSRP}}$ (dB)	IoU*	F1*	Cond. RMSE_h (m)
IDW	16.46	–	–	–
Kriging–Raw	13.68	–	–	–
MLP-RawCoord	12.71	–	–	–
R-Tikh	57.66	0.24	0.39	–
SART	19.68	0.04	0.08	–
NoProtocol	15.70	0.61	0.76	12.4
VO-Compat	5.81	0.66	0.80	5.5
EnvSem-Compat	5.70	0.66	0.80	5.5
Prop	5.24	0.73	0.84	4.8

Note: "–" indicates that a method does not output the corresponding structure or a finite height proxy. Superscript $\star$ denotes the upper envelope over a threshold scan against reference support. The reference-based threshold scan is used for post-fit structural evaluation. Conditional RMSE_h is computed only for truly occupied cells recovered under $\tau^\star$ with finite height proxies. The valid-cell counts are 599/824, 616/824, and 646/824 for NoProtocol, VO-Compat/EnvSem-Compat, and Prop, respectively. Ground-truth-free controlled-scene reference height and full support results are reported separately in Table 1.

Figure 5 and Table 3 test stability under random splits, source-specific link budget biases, measurement noise, and outliers. Table 3 identifies variability as the sample standard deviation across 200 repetitions. All 200 repetitions completed successfully. With common random seeds, Prop improved all four metrics relative to NoProtocol in all 200 repetitions. The paired mean changes were a 13.9 dB reduction in $\mathrm{RMSE}_{\mathrm{RSRP}}$, gains of 0.10 in IoU* and 0.07 in F1*, and a 7.4 m reduction in RMSE_h. Prop achieved a mean prediction error of 5.6 ± 0.8 dB. The known-antenna diagnostic yielded 5.6 ± 0.8 dB. This proximity shows that constrained antenna and bias handling retains near-known-antenna predictive performance under non-cooperative observation.

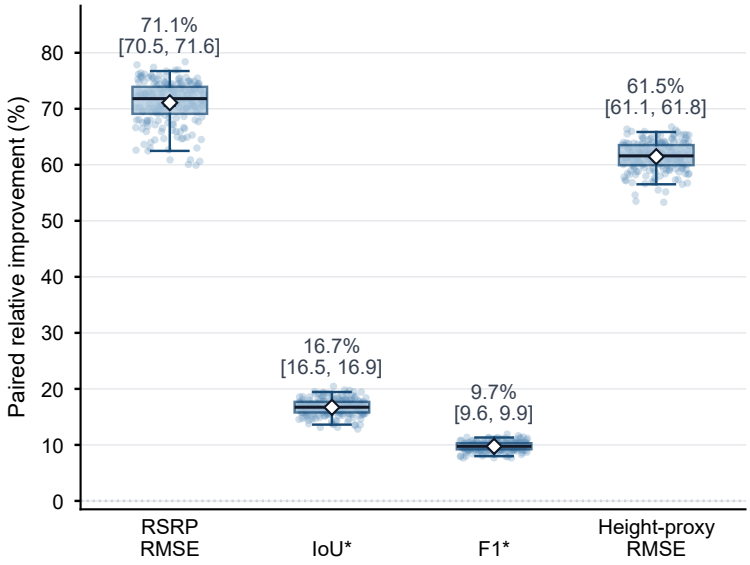


Figure 5. Paired Monte Carlo gains of Prop over NoProtocol. The asterisk denotes scores evaluated at the reference-based IoU-maximizing threshold.

Table 3. Stability across 200 Monte Carlo repetitions (mean $\pm$ standard deviation).

Method	$\mathrm{RMSE}_{\mathrm{RSRP}}$ (dB)	IoU*	F1*	Cond. RMSE_h (m)
IDW	17.3 ± 1.6	–	–	–
RBF-Gauss	19.1 ± 1.7	–	–	–
R-Tikh	53.0 ± 3.8	0.24 ± 0.01	0.38 ± 0.01	–
SART	21.3 ± 1.9	0.05 ± 0.01	0.09 ± 0.01	–
NoProtocol	19.6 ± 2.3	0.61 ± 0.01	0.76 ± 0.01	12.0 ± 0.6
Prop	5.6 ± 0.8	0.72 ± 0.02	0.83 ± 0.01	4.6 ± 0.2

Note: Each repetition varies estimation and holdout road blocks, source-specific biases, measurement noise, and outliers. "–" and $\star$ follow Table 2. Each repetition computes RMSE_h under the same conditions, and its valid-cell count varies with recovered support.

6.5. Reconstruction and Protocol Ablation

Figures 6 and 7 present the reconstructed propagation-equivalent structure and the corresponding held-out prediction diagnostics.

Figure 6 shows spatial reconstructions from all six algorithms for the same non-cooperative RSRP case. Table 2 gives the quantitative results. R-Tikh and SART output only two-dimensional support, and the figure shows that they do not estimate height. NoProtocol overfills the scene after the combined protocol and geometric anchors are removed. VO-Compat and EnvSem-Compat are shown separately but share the same underlying vREM. The residual Kriging head in EnvSem-Compat changes only held-out RSRP predictions. Prop produces support and height proxies that agree more closely with the reference structure than NoProtocol and the compatible reconstructions.

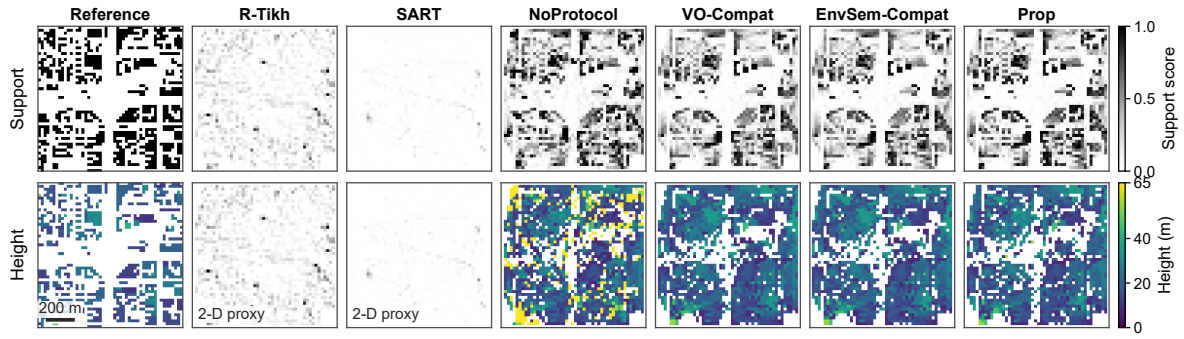


Figure 6. Occupancy support and height proxies of the propagation-equivalent vREM.

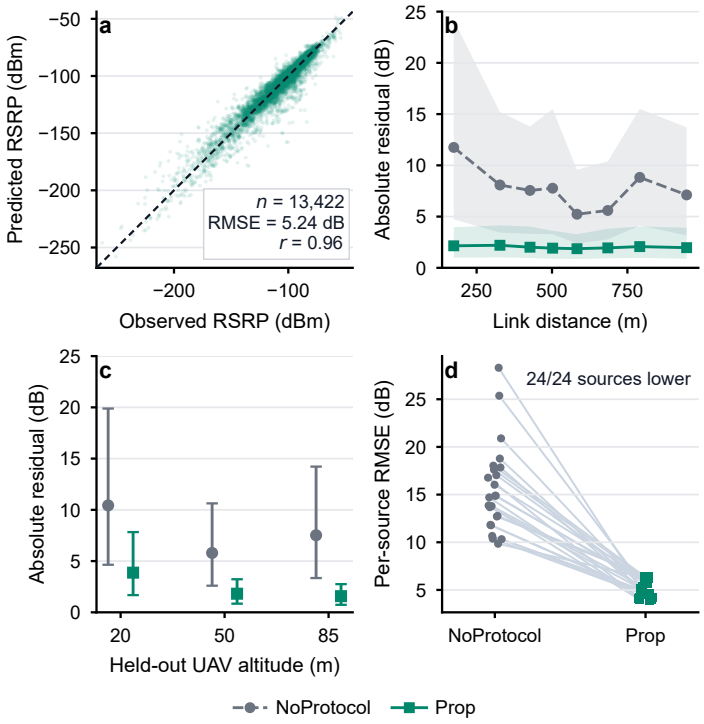


Figure 7. Held-out RSRP prediction diagnostics. (**a**) Predicted versus observed RSRP; (**b**) distance-binned absolute residual; (**c**) altitude-binned absolute residual; (**d**) per-source RMSE.

Figure 7 examines the aggregate performance in Table 2 at the held-out observation level. Prop achieves a prediction–observation Pearson correlation of 0.96. Its absolute residual is lower than NoProtocol across all eight distance bins and three held-out altitudes. Its per-source $\mathrm{RMSE}_{\mathrm{RSRP}}$ also decreases for all 24 cellular sources. The anchoring gains are observed across distance bins, held-out altitudes, and sources. These bins describe the Xi'an scene. Figure 5 and Table 3 report stability across random seeds.

6.6. Robustness and Deployment Diagnostics

To separate robustness from the matched forward/inverse construction, we ran auditable stress tests using the original simulator and estimator. For decoded nominal-power errors of -8 to $+8$ dB, held-out RSRP RMSE stays within 5.20–5.40 dB, with $\mathrm{IoU}^{\star} \approx 0.73$, $\mathrm{F1}^{\star} \approx 0.84$, and conditional height RMSE at 4.8–5.0 m. Changing the nuisance bias bound among ± 4, ± 8, and ± 12 dB yields the same rounded main metrics, and recovered nominal-scene biases achieve 1.48 dB RMSE.

Perturbed attenuation scales and path-loss exponents produce RSRP RMSE ranges of 4.05–6.49 dB and 5.30–5.59 dB, respectively, with structural scores near the nominal result.

Under antenna survey errors of 0–30°, azimuth errors give 4.73–4.88 dB and downtilt errors 4.90–4.98 dB, with structural scores remaining near 0.61/0.76 (IoU/F1). Complete paired sequences and source-level diagnostics are reported in Appendix C. To test source-to-source rather than common-mode power error, the same zero-mean error direction was scaled to standard deviations of 0, 2, 4, 6, 8, and 10 dB. The corresponding RSRP RMSE values were 5.24, 5.26, 5.29, 5.44, 5.71, and 6.17 dB; IoU* values were 0.728, 0.726, 0.725, 0.722, 0.720, and 0.718; and boundary-hit counts were 0, 0, 1, 2, 2, and 8 of 24 sources. Thus, residual source heterogeneity causes gradual degradation rather than an abrupt failure, while boundary hits provide a direct diagnostic when the calibrated operating range is exceeded.

The DPM path-gain maps are used as the sole propagation input, while the building raster is withheld from fitting and used only for post-fit structural evaluation. Because these data contain path gain rather than LTE/NR broadcast fields, this experiment tests deterministic free-space/path-loss normalization and the obstruction inverse under an independent DPM forward process, complementing the real-air interface protocol power evaluation.

Table 4 summarizes the independent DPM results.

Table 4. Independent RadioMapSeer/WinProp-DPM mismatch test. City-level median [interquartile range] over 99 cities.

Method	Holdout RMSE (dB)	AP	IoU	F1	Chamfer (m)
Normalized + bounded offset	6.23 [4.96, 6.97]	0.67 [0.61, 0.74]	0.46 [0.39, 0.52]	0.63 [0.56, 0.68]	6.7 [5.6, 8.3]
Unnormalized + offsets	23.77 [22.72, 24.86]	0.45 [0.35, 0.54]	0.35 [0.25, 0.44]	0.51 [0.40, 0.61]	8.5 [6.6, 12.2]
Free-space only	15.71 [13.02, 18.50]	–	–	–	–
Per-source median	10.93 [10.19, 11.39]	–	–	–	–

Figure 8 illustrates these results. Normalized + bounded offset attains a median holdout RMSE of 6.23 dB, versus 23.77 dB for the unnormalized inverse, 15.71 dB for free space, and 10.93 dB for the per-source median. Relative to the unnormalized inverse, RMSE and building cell AP improve on all 99 cities; the median paired RMSE reduction is 17.95 dB (bootstrap 95% interval 17.82–18.11 dB) and the median AP increase is 0.20 (0.19–0.24).

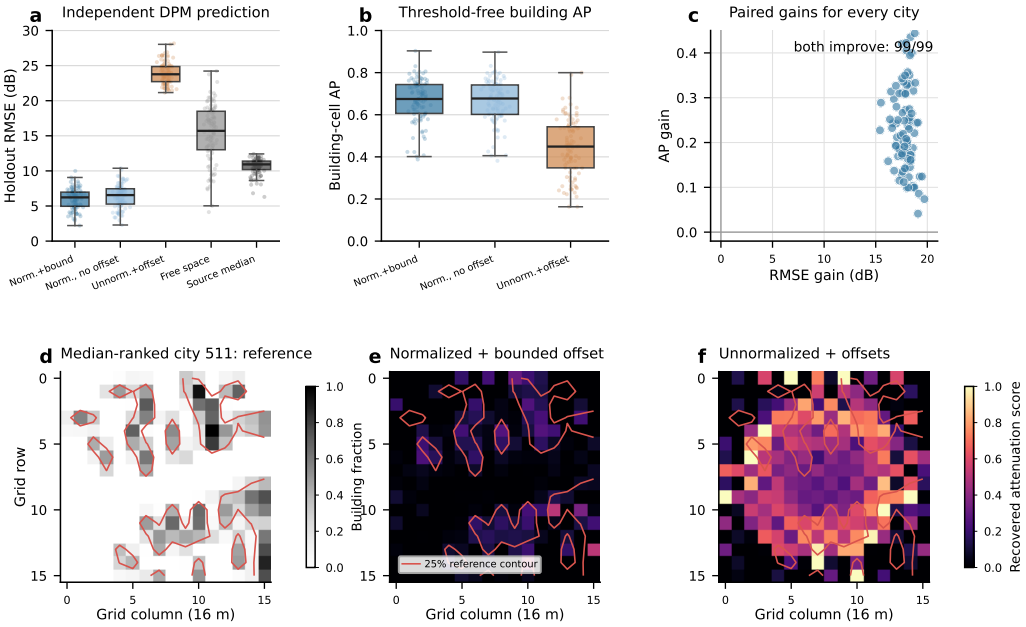


Figure 8. Independent WinProp-DPM mismatch audit. (**a**) Holdout RMSE; (**b**) building-cell AP; (**c**) paired gains; (**d**) city 511 reference map; (**e**) normalized reconstruction with bounded offset; (**f**) unnormalized reconstruction with offsets. Red contours mark a reference building fraction of 0.25.

The RadioMapSeer experiment uses an independent two-dimensional dominant-path propagation process with equal 1.5 m transmit and receive heights to test transfer across city geometries. The forward-only two-ray test further probes the response to an added reflected path.

To add a path-level stress test without changing the inverse solver, we generated a forward-only two-ray ground-reflection residual. The direct path and a ground-image path were combined coherently at $|\Gamma| = 0.30$, 0.55, and 0.80, while the inverse model remained the original Beer–Lambert height-class estimator. The full scene and train/holdout construction were retained. As $|\Gamma|$ increased, the injected residual standard deviation increased from 1.81 to 5.34 dB, and Prop's held-out RSRP RMSE increased from 5.58 to 7.63 dB. In contrast, oracle IoU remained 0.73–0.72 and conditional height RMSE remained 4.7–4.8 m; NoProtocol remained at 15.78–16.53 dB RSRP RMSE and 12.2–12.7 m conditional height RMSE (Table 5). The Prop–NoProtocol prediction, structural, and height advantages therefore remain under the added path-level residual.

Table 5. Forward two-ray ground-reflection mismatch.

$\|\Gamma\|$	Residual SD (dB)	Prop RMSE (dB)	Prop IoU*	Prop Height RMSE (m)	NoProtocol RMSE (dB)	NoProtocol Height RMSE (m)
0.00 (nominal)	0.00	5.24	0.73	4.8	15.70	12.4
0.30	1.81	5.58	0.73	4.7	15.78	12.5
0.55	3.44	6.34	0.73	4.7	16.03	12.2
0.80	5.34	7.63	0.72	4.8	16.53	12.7

All settings use the same full scene, random seed, estimator and disjoint holdout construction. The two-ray term is present only in forward generation; the inverse does not observe $|\Gamma|$. Superscript $\star$ denotes reference-based oracle IoU.

Table 6 summarizes the sensitivity tests.

Table 6. Robustness and sensitivity tests for the proposed estimator. Superscript $\star$ denotes scores at the reference-based IoU-maximizing threshold.

Setting	RMSE (dB)	IoU*	F1*	Height RMSE (m)
Nominal, 11 classes, 20 m	5.24	0.73	0.84	4.8
Nominal-power error $-8/+8$ dB	5.31/5.40	0.73/0.73	0.84/0.84	4.8/5.0
Bias bound $\pm4/\pm8/\pm12$ dB	5.24/5.24/5.24	0.73/0.73/0.73	0.84/0.84/0.84	4.8/4.8/4.8
True $\eta = 1.0/2.0$ dB/m	4.05/6.49	0.73/0.73	0.84/0.84	4.8/4.7
True path-loss exponent 1.8/2.5	5.59/5.30	0.73/0.73	0.84/0.84	4.9/4.8
Survey azimuth error	4.73/4.73/4.72	0.61/0.61/0.61	0.76/0.76/0.76	5.3/5.0/5.0
0/5/10/20/30° (matched)	4.76/4.88	0.61/0.61	0.76/0.76	5.2/5.3
Survey downtilt error	4.73/4.90/4.95	0.61/0.61/0.61	0.76/0.76/0.76	5.3/5.3/5.3
0/5/10/20/30° (matched)	4.98/4.98	0.61/0.61	0.76/0.76	5.5/5.5

The C2 candidate source scan was complemented by a fixed-budget scan of independent physical sites and crossing angles. Relative to one site, two independent sites reduce RSRP RMSE from 12.46 to 10.83 dB, increase IoU* from 0.241 to 0.319, and reduce conditional height RMSE from 10.4 to 8.2 m; the 60°–120° crossing-angle region provides the most effective joint geometry. C3 uses 7/5/3/2/1 altitude layers with a fixed holdout altitude set. From one to seven layers, RSRP RMSE changes from 11.04 to 5.64 dB and conditional height RMSE from 12.5 to 5.6 m; the two-layer case establishes the first vertical contrast, while the seven-layer setting provides the strongest combined prediction and height result.

The matched-budget scans connect the geometric diagnostics to finite-sample reconstruction. As the C1 window contracts from 60° to 30°, held-out RSRP RMSE increases

from 6.64 to 11.15 dB and IoU* decreases from 0.69 to 0.54. Increasing C3 from two to seven altitude layers reduces RSRP RMSE from 11.36 to 5.64 dB and conditional height RMSE from 7.7 to 5.6 m. Together with the C2 site and angle comparisons above, these scans provide an empirical link between C1–C3 geometry and reconstruction quality.

In the Xi'an scene, $\Theta_{3\mathrm{dB}} = 65^{\circ}$, the median occupied-building height is 18.6 m, 22/24 sectors meet the C1 azimuth target, occupied multi-source coverage is 0.88, and seven layers span 95 m (Appendix A). The most effective crossing-angle band is 60°–120°, and three or more altitude layers strengthen height identification. These diagnostics provide measurable deployment checks for a new route.

For deployment-oriented thresholding, the oracle scan is reported only as an upper-envelope diagnostic. Table 7 gives the complete nominal audit. For Prop, the 256-bin Otsu and oracle IoU values both round to 0.73; under the same Otsu protocol, NoProtocol yields 0.61, so Prop retains a rounded IoU advantage of 0.12 without using reference support to select the threshold.

Table 7. Nominal oracle and ground-truth-free occupancy threshold audit.

Method	Rule	τ	IoU
Prop	Oracle	0.3750	0.73
Prop	Otsu (unsupervised)	0.4043	0.73
Prop	Fixed midpoint	0.5000	0.67
NoProtocol	Oracle	0.3750	0.61
NoProtocol	Otsu (unsupervised)	0.4043	0.61
NoProtocol	Fixed midpoint	0.5000	0.56

Note: Oracle thresholds use reference support and are diagnostic only. Fixed and Otsu thresholds do not read reference support during selection.

For the grid-resolution audit, observations are generated once on a fixed 5 m forward reference grid and the identical 40,000/4869 train/holdout tuples are inverted at 10, 20, and 40 m. The corresponding RSRP RMSE values are 7.24/7.86/9.30 dB and the oracle IoU values are 0.40/0.45/0.43. The 20 m grid provides the highest support score at intermediate memory, while 10 m provides lower prediction error at larger memory and 40 m provides lower memory at coarser spatial representation. These measurements identify 20 m as the balanced prediction–support–resource resolution for the 1000 m scene.

To isolate class discretization from height-range effects, the 8/11/16-class audit uses the identical 5–65 m envelope and the same 123,395/13,422 train/holdout tuples. The three settings give a held-out RSRP RMSE of 6.99/5.24/5.23 dB, an IoU* of 0.66/0.73/0.68, and a conditional height RMSE of 4.7/4.8/4.4 m. The 11-class dictionary retains the highest support score, while the 16-class dictionary gives comparable prediction and height accuracy at a larger variable count. This comparison quantifies the accuracy–resolution trade-off across class settings.

The nominal full run was profiled on an Intel Core i9-11900KB workstation (8 cores/16 threads), with the sparse matrix construction requiring 12.33 s and the three-update Prop inversion requiring 65.61 s. The complete measurement-to-map update chain therefore takes 77.94 s with a 1.648 GB peak working set. Sparse ray–cell products and source-wise updates can be batched on a GPU, reused across rolling windows, and warm-started from the previous map, providing the computational basis for online vREM updates on an airborne or edge-computing platform.

The controlled simulation evaluates complete multi-altitude 2.5D inversion. We next test whether a real ground-based passive receiver chain can implement protocol anchoring and multi-source evidence formation.

6.7. Real-Road Passive Mapping

The field experiment tests a direct deployment claim: a ground-based receiver chain can form a spatially stable ground-level vREM from downlink IQ, decodable protocol fields, sites, GPS, and road geometry. We examine relative-power mapping, reproducibility across interleaved acquisition folds, and correspondence with block boundaries. These tests evaluate protocol recovery and multi-source spatial evidence formation in a real receiver chain. Figure 9 summarizes the acquisition geometry.

The field platform used a YunSDR GRFM-7400 software-defined radio system running software v1.2.1.9. Its FMC9371 RF module contains an AD9371 transceiver. The equipment specification lists up to 4 × 4 channels, a 300 MHz–6 GHz RF range, and a maximum receive bandwidth of 100 MHz; these are the hardware capabilities rather than the settings used in this experiment. The installed antenna was a 100 MHz–6 GHz wideband omnidirectional antenna. The experiment enabled two RF channels and cycled through five LTE center frequencies (939, 1815, 1895, 2017.5, and 2624.6 MHz) and one 5G NR center frequency (2524.95 MHz). It recorded 30.72 Msps complex IQ with a 500 ms timestamp guard. Formal route metadata record a fixed 60 dB software-commanded receiver gain value at all six frequencies, with the external LNA enabled. A GPSDO supplied PPS timing, and the GPS antenna supplied single-point position fixes for spatial aggregation. The platform produced 1518 valid multi-source observations covering 82 LTE/NR sources and 41 physical sites. Figure 9 shows the trajectory, multi-source sites, observation rays, and mapping region.

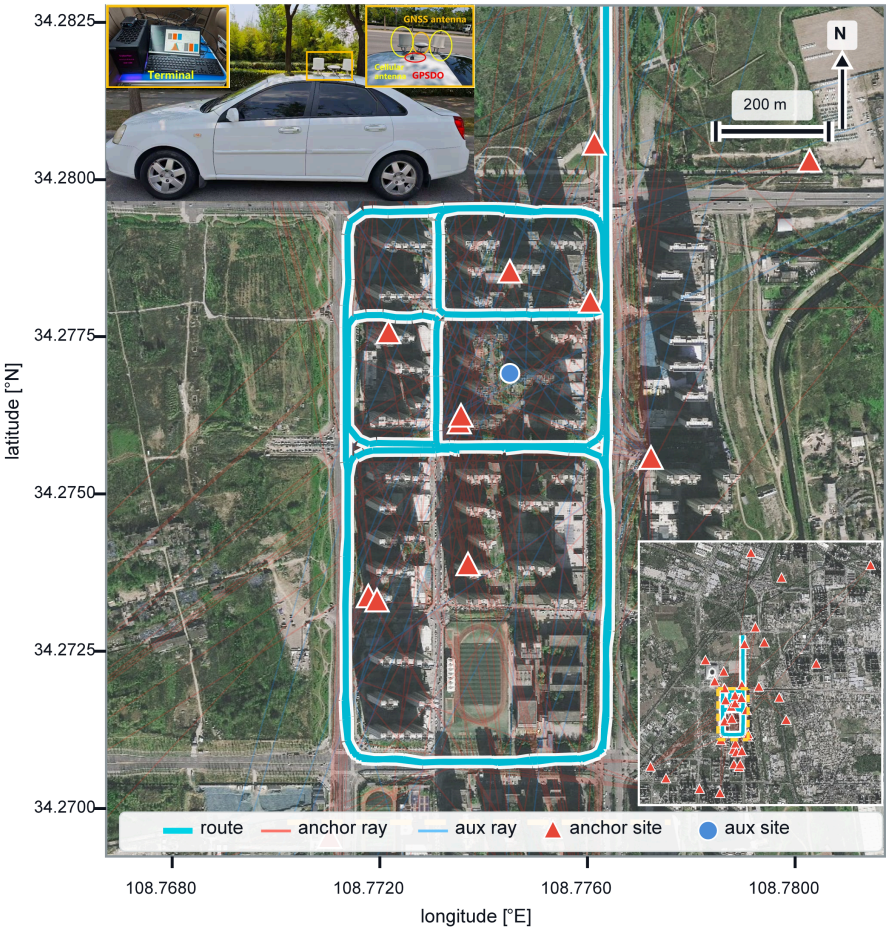


Figure 9. Field acquisition geometry. The Amap satellite image is used only as a display basemap [52].

We additionally tested whether the decoded LTE power configuration tracks an independently constructed received scale in the commercial–air interface data. The analysis

used 63 LTE cells from 28 physical sites. The independent scale combined coherent CRS digital power with deterministic free-space loss and did not use `referenceSignalPower`. We centered both variables within six receiver-front-end categories and used physical-site clusters for uncertainty estimation. Each 1 dB increase in configured CRS EPRE corresponded to a 0.819 dB increase in the independent received-scale proxy. The Huber slope was 0.883, and leave-one-site slopes ranged from 0.568 to 1.089. The site-clustered bootstrap 95% interval was 0.030–1.539 dB/dB. A within-front-end stratified permutation test rejected zero association ($p = 0.00665$). This field result links the broadcast power configuration to a measurable received scale without requiring absolute receiver chain calibration.

Before multi-source fusion, 154 fixed-gain 1815 MHz segments test whether relative-power observations support spatial mapping. The Gaussian-kernel bandwidth is fixed at the 35 m value used by the revisit protocol. Observations separated by less than 120 s are excluded at each evaluation point. Across 116 points with time-separated revisits, the ground-level relative-power field achieves a Pearson correlation of 0.85 and an RMSE of 3.41 dB. The global-median predictor yields an RMSE of 7.27 dB. Spatial mapping reduces error by 53.0% relative to this baseline. As a spatial-scale check, the median absolute power difference is 2.86 dB within 35 m and 9.04 dB beyond 150 m. This benchmark directly evaluates spatial prediction of a relative-power field on a fixed digital scale.

The ground-level vREM uses a source-normalization, segment-aggregation, and spatial-mapping pipeline. Robust within-source normalization is applied to protocol path-loss residuals, relative power, and signal quality; multi-band evidence is mapped onto a 20 m grid, and availability-normalized fusion produces a continuous ground-level radio evidence field. Corrected GPS trajectories and non-building road-semantic context define the candidate domain; uniform thresholding and connectivity processing extract block-scale candidate regions for OSM boundary evaluation. The same pipeline is used for every band and temporal fold.

To quantify how much structure is supplied by the cellular observations, we use a fixed-area, same-domain audit. Context-only, Radio-only, and Joint share the same 796-cell road/trajectory-conditioned candidate domain and the same 274-cell output budget (10% of the 2739 measurement-supported cells and 34.4% of the common domain). Context-only uses road and trajectory context, Radio-only uses protocol-normalized power, visibility, signal quality, and temporal consistency, and Joint fuses the two rank fields. OSM building geometry [53] is read only after all three masks are frozen. The resulting fixed-budget comparison is reported in Table 8.

Table 8. Fixed-area comparison of cellular and contextual evidence in the ground-level field map.

Method	AUC	Precision	Recall	IoU	P/C (20 m)	P/C (40 m)	Chamfer (m)
Context-only	0.584	0.226	0.320	0.153	0.628/0.560	0.805/0.798	38.18
Radio-only	0.583	0.266	0.376	0.185	0.660/0.487	0.865/0.560	41.54
Joint	0.634	0.325	0.459	0.235	0.744/0.627	0.890/0.756	30.98

At the same 10% output budget, the OSM-blind radio texture gives building-ranking AUC 0.716 on 1999 measurement-supported cells; a time/source-preserving cyclic shift null has a 95th-percentile AUC of 0.678 ($p = 0.03247$). The fixed-area audit and the OSM-blind test together show that cellular observations supply independent building-related spatial information, while the joint rank field concentrates that information into the most accurate candidate regions.

Table 9 summarizes the field results.

Table 9. Field results for the ground-level vREM.

Metric	Result
Time-separated mapping points	116
Relative-power prediction Pearson/RMSE	0.85/3.41 dB
Global-median baseline RMSE	7.27 dB
Interleaved-fold mask Jaccard	0.97
Boundary precision/coverage (20 m)	0.58/0.66
Boundary precision/coverage (40 m)	0.80/0.81
Symmetric boundary Chamfer distance	33.7 m
OSM-blind radio-texture AUC	0.716 ($p = 0.03247$)
Fixed 10% budget Joint AUC/Precision/Recall/IoU	0.634/0.325/0.459/0.235
Fixed 10% budget Joint boundary P/C (20/40 m)	0.744/0.627; 0.890/0.756

Figure 10 shows the continuous radio score and its spatial correspondence with OSM building boundaries. At a 40 m boundary-matching tolerance, candidate-boundary precision is 0.80 and reference-boundary coverage is 0.81, with a symmetric boundary Chamfer distance of 33.7 m. Together with the 53.0% relative-power RMSE reduction and 0.97 interfold Jaccard, these results establish stable block-scale ground-level propagation evidence for macro-scale route screening and connectivity assessment.

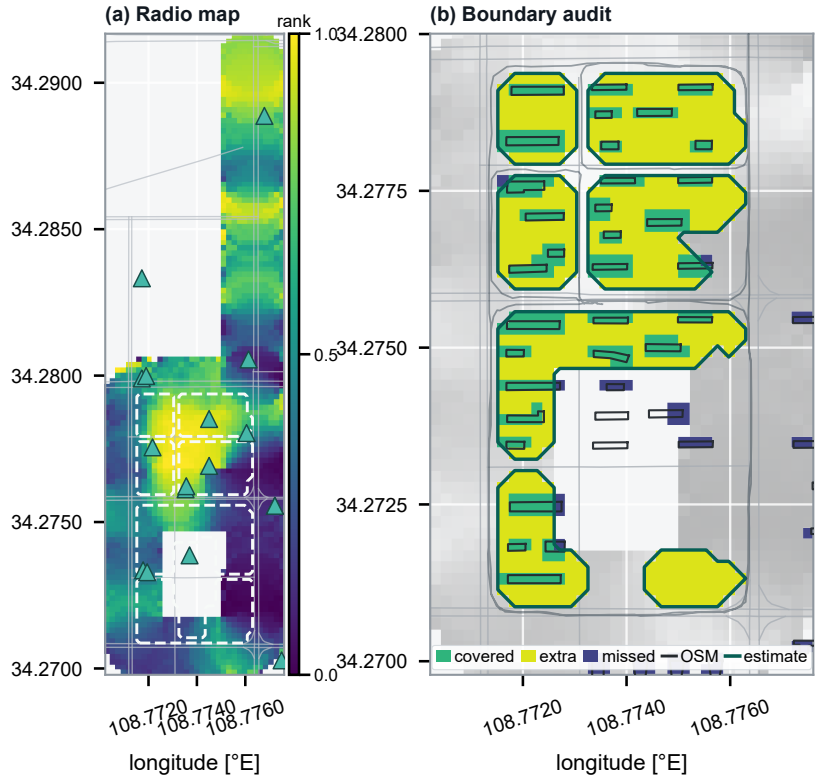


Figure 10. Correspondence between the ground-level vREM and OSM boundaries. Triangles indicate cellular sites; white dashed outlines mark the extracted candidate regions.

7. Conclusions

We proposed a protocol-anchored, propagation-equivalent vREM that uses decoded nominal power, path-loss normalization, coarse site bearings, and multi-altitude trajectories to separate environmental structure from transmitter-side uncertainty. In the Xi'an simulation, the method reduced held-out RSRP RMSE from 15.70 to 5.24 dB and produced stable gains across repeated trials. On the fixed RadioMapSeer split [54], deterministic

normalization reduced median DPM holdout error from 23.77 to 6.23 dB and improved building cell AP in all 99 cities, supporting transfer across urban geometries.

The real-road experiment verified protocol recovery and stable ground-level radio evidence formation from operational LTE/NR downlinks without building geometry during inference, reducing relative-power RMSE by 53.0% over the global-median baseline. The full 2.5D height result remains simulation-based, and recovered height is a propagation-equivalent proxy; the field result supports block-scale structural correspondence. Together, these findings position vREM as a complementary evidence layer for low-altitude route screening and communication risk assessment.

Author Contributions: Conceptualization, Z.B.; methodology, Z.W.; software, H.L. (Hu Lu); validation, C.Z. and H.L. (Hongcheng Li); formal analysis, H.L. (Hu Lu); investigation, L.W.; resources, F.H.; data curation, X.H.; writing—original draft preparation, Z.B.; writing—review and editing, H.L. (Hu Lu); visualization, Z.W.; supervision, H.L. (Hongcheng Li); project administration, L.W.; funding acquisition, H.L. (Hu Lu). All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Research on Human–UAV Cooperative Visual Language Navigation Technology in Counter-Attacking Scenarios of National Key Laboratory of Unmanned Aerial Vehicle Technology, grant number WRFX202507. The APC was funded by National Key Laboratory of Unmanned Aerial Vehicle Technology.

Data Availability Statement: The reproducibility package is publicly available at https://github.com/scorpio-bza/vrem-reproducibility-package (accessed on 8 September 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

3GPP	Third Generation Partnership Project
ASN.1	Abstract Syntax Notation One
AUC	Area under the curve
CKM	Channel knowledge map
CRC	Cyclic redundancy check
EPRE	Energy per resource element
FDD	Frequency-division duplex
FIM	Fisher information matrix
FISTA	Fast iterative shrinkage-thresholding algorithm
GPR	Gaussian-process regression
GPS	Global Positioning System
IDW	Inverse-distance weighting
IQ	In-phase and quadrature
ISAC	Integrated sensing and communication
LSQR	Least-squares QR algorithm
LTE	Long-Term Evolution
LTE-CRS	LTE cell-specific reference signal
MDT	Minimization of drive tests
MIB	Master information block
MLP	Multilayer perceptron
NLOS	Non-line-of-sight
NR	New Radio

OSM	OpenStreetMap
PBCH	Physical broadcast channel
RBF	Radial basis function
REM	Radio environment map
RMSE	Root mean square error
RRC	Radio resource control
RSRP	Reference signal received power
RSS	Received signal strength
RTI	Radio tomographic imaging
SART	Simultaneous algebraic reconstruction technique
SIB	System information block
SLAM	Simultaneous localization and mapping
SSB	Synchronization signal/PBCH block
SS-RSRP	Synchronization signal reference signal received power
TDD	Time-division duplex
UAV	Unmanned aerial vehicle
VO	Virtual obstacle
vREM	Virtual radio environment map

Appendix A. Geometric Coverage Diagnostics

Table A1 reports the C1–C3 diagnostics for the Xi'an experiment, as defined in Section 4.

Table A1. Observation geometry coverage diagnostics.

Cond.	Diagnostic	Result
C1	Sector azimuth-span target of approximately $2\Theta_{3dB}$	22/24 meet target; minimum/median spans 128.7°/354.6°
C2	Occupied-cell multi-source coverage and maximum source separation relative to 65°	Multi-source coverage 0.88; angular-target rate 0.78; median source count 18
C3	Multiple altitude layers with span comparable to the median occupied-building height	7 layers spanning 95.0 m; median building height 18.6 m; target met

Appendix B. Estimator and Baseline Settings

Prop uses the height classes defined in Section 6.1, $\lambda_1 = 2.0$, and 180 nonnegative FISTA updates per environmental block. Environmental and bias blocks alternate three times with $\delta_b \in [-8, 8]$ dB. Three outer antenna updates refine azimuth and downtilt within the constrained search domain.

All baselines use configurations fixed before formal holdout evaluation. Kriging–Raw uses 16 neighbors with a 60 m length scale; MLP-RawCoord is a two-layer ReLU network trained with Adam and early stopping; VO-Compat and EnvSem-Compat use $\lambda = 0.5$ with three bias updates; IDW uses power 2; RBF-Gauss uses a 140 m length scale; R-Tikh uses ridge regularization; and SART uses eight iterations. Complete hyperparameter settings are provided in the public code repository.

Appendix C. Sensitivity Analyses

This appendix reports the complete sensitivity tables supporting Section 6.6; all cases use matched train/holdout budgets and fixed evaluation rules.

Table A2 reports nominal-power sensitivity.

Table A2. Decoded nominal-power sensitivity. Superscript ⋆ denotes reference-based oracle IoU.

Power Error (dB)	RMSE (dB)	IoU⋆	Height RMSE (m)
−8	5.31	0.73	4.8
−4	5.20	0.73	4.8
0	5.24	0.73	4.8
+4	5.30	0.73	4.8
+8	5.40	0.73	5.0

The attenuation scale ($\eta_0 = 1.00$–$2.00\,\mathrm{dB/m}$) and sparsity ($\lambda_1 = 0.1$–10) sweeps are flat across settings, supporting $\eta_0 = 1.5\,\mathrm{dB/m}$ and $\lambda_1 = 2$ as fixed reconstruction settings.

The resolution audit inverts identical observations at 10, 20, and 40 m; the 20 m grid provides the balanced prediction–support–resource trade-off for the 1000 m scene.

Table A3 reports attenuation-scale and sparsity sensitivity.

Table A3. Attenuation scale and sparsity sensitivity. Superscript ⋆ denotes reference-based oracle IoU.

Audit	Tested Settings	Train/Test	RMSE Range	IoU⋆ Range	Height Range
Inverse η_0 (dB/m)	1.00–2.00 (5)	41,188/4890	4.73–4.73	0.61–0.61	5.1–5.3 m
λ_1	0.1–10 (6)	41,188/4890	4.73–4.73	0.61–0.61	5.2–5.3 m

Table A4 reports the grid-resolution audit.

Table A4. Grid resolution audit under matched observations. Superscript ⋆ denotes reference-based oracle IoU.

Grid (m)	Variables	Peak Memory (MB)	RMSE (dB)	IoU⋆	Height (m)
10	110,000	1031	7.24	0.40	7.2
20	27,500	632	7.86	0.45	6.9
40	6875	433	9.30	0.43	6.6

Table A5 reports additional sensitivity results.

Table A5. Additional source heterogeneity, discretization, and reduced budget sensitivity results. Superscript ⋆ denotes reference-based oracle IoU.

Case	Train Tuples	Grid (m)	RMSE (dB)	IoU⋆	Height (m)
Source-power error SD 0/2/4/6/8/10 dB	123,395	20	5.24/5.26/5.29/5.44/ 5.71/6.17	0.728/0.726/0.725/0.722/ 0.720/0.718	–
8/11/16 height classes (matched)	123,395	20	6.99/5.24/5.23	0.66/0.73/0.68	4.7/4.8/4.4
Survey azimuth error 5/10/20/30° (reduced)	41,188	20	4.73/4.72/4.76/4.88	0.61/0.61/0.61/0.61	5.0/5.0/5.2/5.3
Survey downtilt error 5/10/20/30° (reduced)	41,188	20	4.90/4.95/4.98/4.98	0.61/0.61/0.61/0.61	5.3/5.3/5.5/5.5

References

1. Lin, X.; Yajnanarayana, V.; Muruganathan, S.D.; Gao, S.; Asplund, H.; Määttänen, H.-L.; Bergström, M.; Euler, S.; Wang, Y.-P.E. The sky is not the limit: LTE for unmanned aerial vehicles. *IEEE Commun. Mag.* **2018**, *56*, 204–210. [CrossRef]
2. Galkin, B.; Kibiłda, J.; DaSilva, L.A. Coverage analysis for low-altitude UAV networks in urban environments. In Proceedings of the IEEE Global Communications Conference (GLOBECOM), Singapore, 4–8 December 2017; pp. 1–6. [CrossRef]
3. Cai, X.; Izydorczyk, T.; Rodriguez-Pineiro, J.; Kovacs, I.Z.; Wigard, J.; Tavares, F.M.L.; Mogensen, P.E. Empirical low-altitude air-to-ground spatial channel characterization for cellular networks connectivity. *IEEE J. Sel. Areas Commun.* **2021**, *39*, 2975–2991. [CrossRef]

4. Liu, F.; Cui, Y.; Masouros, C.; Xu, J.; Han, T.X.; Eldar, Y.C.; Buzzi, S. Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond. *IEEE J. Sel. Areas Commun.* **2022**, *40*, 1728–1767. [CrossRef]
5. Wilson, J.; Patwari, N. Radio tomographic imaging with wireless networks. *IEEE Trans. Mob. Comput.* **2010**, *9*, 621–632. [CrossRef]
6. Zeng, Y.; Chen, J.; Xu, J.; Wu, D.; Xu, X.; Jin, S.; Gao, X.; Gesbert, D.; Cui, S.; Zhang, R. A tutorial on environment-aware communications via channel knowledge map for 6G. *IEEE Commun. Surv. Tutor.* **2024**, *26*, 1478–1519. [CrossRef]
7. Kassas, Z.M.; Khalife, J.; Shamaei, K.; Morales, J. I hear, therefore I know where I am: Compensating for GNSS limitations with cellular signals. *IEEE Signal Process. Mag.* **2017**, *34*, 111–124. [CrossRef]
8. Sirola, N. Closed-form algorithms in mobile positioning: Myths and misconceptions. In Proceedings of the 7th Workshop Positioning, Navigation and Communication, Dresden, Germany, 11–12 March 2010; pp. 38–44. [CrossRef]
9. Levie, R.; Yapar, C.; Kutyniok, G.; Caire, G. RadioUNet: Fast radio map estimation with convolutional neural networks. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 4001–4015. [CrossRef]
10. Teganya, Y.; Romero, D. Deep completion autoencoders for radio map estimation. *IEEE Trans. Wirel. Commun.* **2022**, *21*, 1710–1724. [CrossRef]
11. Zeng, Y.; Xu, X. Toward environment-aware 6G communications via channel knowledge map. *IEEE Wirel. Commun.* **2021**, *28*, 84–91. [CrossRef]
12. Chen, J.; Esrafilian, O.; Gesbert, D.; Mitra, U. Efficient algorithms for air-to-ground channel reconstruction in UAV-aided communications. In Proceedings of the IEEE Globecom Workshops (GC Wkshps), Singapore, 4–8 December 2017; pp. 1–6. [CrossRef]
13. Liu, W.; Chen, J. UAV-aided radio map construction exploiting environment semantics. *IEEE Trans. Wirel. Commun.* **2023**, *22*, 6341–6355. [CrossRef]
14. Chen, W.; Chen, J. Diffraction and scattering aware radio map and environment reconstruction using geometry model-assisted deep learning. *IEEE Trans. Wirel. Commun.* **2024**, *23*, 19804–19819. [CrossRef]
15. Song, Y.; Zeng, Y.; Yang, Y.; Ren, Z.; Cheng, G.; Xu, X.; Xu, J.; Jin, S.; Zhang, R. An overview of cellular ISAC for low-altitude UAV: New opportunities and challenges. *IEEE Commun. Mag.* **2025**, *63*, 88–95. [CrossRef]
16. Jiang, Y.; Li, X.; Zhu, G.; Li, H.; Deng, J.; Han, K.; Shen, C.; Shi, Q.; Zhang, R. Integrated sensing and communication for low-altitude economy: Opportunities and challenges. *IEEE Commun. Mag.* **2025**, *63*, 72–78. [CrossRef]
17. Samczynski, P.; Abratkiewicz, K.; Plotka, M.; Zielinski, T.P.; Wszolek, J.; Hausman, S.; Korbel, P.; Ksiezyk, A. 5G network-based passive radar. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 5108209. [CrossRef]
18. Maksymiuk, R.; Abratkiewicz, K.; Samczynski, P.; Plotka, M. Rényi entropy-based adaptive integration method for 5G-based passive radar drone detection. *Remote Sens.* **2022**, *14*, 6146. [CrossRef]
19. Taylor, A.; Poullin, D. Experimental UAV detection using 4G-LTE-based passive radar. In Proceedings of the IEEE International Radar Conference (RADAR), Sydney, Australia, 6–10 November 2023; pp. 1–6. [CrossRef]
20. Shamaei, K.; Khalife, J.; Kassas, Z.M. Exploiting LTE signals for navigation: Theory to implementation. *IEEE Trans. Wirel. Commun.* **2018**, *17*, 2173–2189. [CrossRef]
21. Vakalis, S.; Mghabghab, S.; Nanzer, J.A. Passive non-cooperative millimeter-wave imaging using 5G signals of opportunity. In Proceedings of the IEEE MTT-S International Microwave Symposium (IMS), Atlanta, GA, USA, 7–25 June 2021; pp. 549–552. [CrossRef]
22. Yiğitler, H.; Jäntti, R.; Kaltiokallio, O.; Patwari, N. Detector-based radio tomographic imaging. *IEEE Trans. Mob. Comput.* **2018**, *17*, 58–71. [CrossRef]
23. Wang, Z.; Guo, X.; Wang, G. Exploring the Laplace prior in radio tomographic imaging with sparse Bayesian learning towards the robustness to multipath fading. *Sensors* **2019**, *19*, 5126. [CrossRef] [PubMed]
24. Wilson, J.; Patwari, N. A fade-level skew-Laplace signal-strength model for device-free localization with wireless networks. *IEEE Trans. Mob. Comput.* **2012**, *11*, 947–958. [CrossRef]
25. Lee, D.; Giannakis, G.B. A variational Bayes approach to adaptive radio tomography. *IEEE Trans. Signal Process.* **2020**, *68*, 3779–3792. [CrossRef]
26. *GPP TS 37.320*; Radio Measurement Collection for Minimization of Drive Tests (MDT); Overall Description; Stage 2, Version 17.6.0. Available online: https://www.3gpp.org/ftp/Specs/archive/37_series/37.320/37320-h60.zip (accessed on 17 June 2025).
27. Enami, R.; Gupta, S.; Rajan, D.; Camp, J. LAIK: Location-specific analysis to infer key performance indicators. *IEEE Trans. Veh. Technol.* **2021**, *70*, 4406–4418. [CrossRef]
28. Li, H.; Li, P.; Cheng, G.; Xu, J.; Chen, J.; Zeng, Y. Channel knowledge map (CKM)-assisted multi-UAV wireless network: CKM construction and UAV placement. *J. Commun. Inf. Netw.* **2023**, *8*, 256–270. [CrossRef]
29. Qiu, Y.; Chen, X.; Mao, K.; Ye, X.; Li, H.; Ali, F.; Huang, Y.; Zhu, Q. Channel knowledge map construction based on a UAV-assisted channel measurement system. *Drones* **2024**, *8*, 191. [CrossRef]

30. Li, Y.; Li, Z.; Gao, Z.; Chen, T. Geo2SigMap: High-fidelity RF signal mapping using geographic databases. In Proceedings of the IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN), Washington, DC, USA, 13–16 May 2024; pp. 277–285. [CrossRef]
31. Yapar, C.; Jaensch, F.; Levie, R.; Kutyniok, G.; Caire, G. Overview of the first pathloss radio map prediction challenge. *IEEE Open J. Signal Process.* **2024**, *5*, 948–963. [CrossRef]
32. Wang, J.; Zhu, Q.; Lin, Z.; Chen, J.; Ding, G.; Wu, Q.; Gu, G.; Gao, Q. Sparse Bayesian learning-based hierarchical construction for 3D radio environment maps incorporating channel shadowing. *IEEE Trans. Wirel. Commun.* **2024**, *23*, 14560–14574. [CrossRef]
33. Chen, W.; Chen, J. Blockage-aware radio map construction via exploiting the diffraction and obstruction structure. In Proceedings of the IEEE Global Communications Conference (GLOBECOM), Kuala Lumpur, Malaysia, 4–8 December 2023; pp. 1920–1925. [CrossRef]
34. Wang, X.; Pan, Y.; Cheng, N.; Yapar, C.; Sun, R.; Yin, Z.; Zhou, C.; Xu, W.; Zhang, Y.; Zhang, J.; et al. A tutorial on learning-based radio map construction: Data, paradigms, and physics-awareness. *arXiv* **2026**, arXiv:2603.17499.
35. Luo, Y.; Wang, Y.; Chen, H.; Wu, C.; Lyu, X.; Zhou, J.; Ma, J.; Zhang, F.; Zhou, B. FERMI: Flexible radio mapping with a hybrid propagation model and scalable autonomous data collection. *arXiv* **2025**, arXiv:2504.14862.
36. Lin, Y.; Xiao, Z.; Wu, H.; Qiao, X.; Wan, F.; Zhang, T. Channel knowledge map enabled low-complexity dynamic radio environment reconstruction. *arXiv* **2026**, arXiv:2608.17322.
37. Fang, Z.; Liu, J.; Liu, K.; Zhang, J.; Wang, Y.; Chen, K. RadioFormer3D: Weakly supervised 3D radio map estimation in low-altitude airspace via generative modeling. *arXiv* **2026**, arXiv:2605.29538.
38. *GPP TS 36.214*; Evolved Universal Terrestrial Radio Access (E-UTRA). Physical Layer; Measurements, Version 16.2.0. Available online: https://www.3gpp.org/ftp/Specs/archive/36_series/36.214/36214-g20.zip (accessed on 16 March 2021).
39. *GPP TS 38.215*; NR; Physical Layer Measurements. Version 18.0.0. Available online: https://www.3gpp.org/ftp/Specs/archive/38_series/38.215/38215-i00.zip (accessed on 18 September 2023).
40. *GPP TS 36.331*; Evolved Universal Terrestrial Radio Access (E-UTRA). Radio Resource Control (RRC); Protocol Specification. Version 16.7.0. Available online: https://www.3gpp.org/ftp/Specs/archive/36_series/36.331/36331-g70.zip (accessed on 16 December 2021).
41. *GPP TS 38.331*; NR; Radio Resource Control (RRC) Protocol Specification. Version 18.0.0. Available online: https://www.3gpp.org/ftp/Specs/archive/38_series/38.331/38331-i00.zip (accessed on 18 December 2023).
42. *GPP TS 38.213*; NR; Physical Layer Procedures for Control. Version 17.9.0. Available online: https://www.3gpp.org/ftp/Specs/archive/38_series/38.213/38213-h90.zip (accessed on 17 March 2024).
43. *GPP TR 38.901*; Study on Channel Model for Frequencies from 0.5 to 100 GHz. Version 17.1.0. Available online: https://www.3gpp.org/ftp/Specs/archive/38_series/38.901/38901-h10.zip (accessed on 17 December 2023).
44. Sun, S.; Rappaport, T.S.; Thomas, T.A.; Ghosh, A.; Nguyen, H.C.; Kovacs, I.Z.; Rodriguez, I.; Koymen, O.; Partyka, A. Investigation of prediction accuracy, sensitivity, and parameter stability of large-scale propagation path loss models for 5G wireless communications. *IEEE Trans. Veh. Technol.* **2016**, *65*, 2843–2860. [CrossRef]
45. Gudmundson, M. Correlation model for shadow fading in mobile radio systems. *Electron. Lett.* **1991**, *27*, 2145–2146. [CrossRef]
46. *GPP TR 36.873*; Study on 3D Channel Model for LTE. Version 12.7.0. Available online: https://www.3gpp.org/ftp/Specs/archive/36_series/36.873/36873-c70.zip (accessed on 12 December 2017).
47. Cherif, N.; Nadeem, Q.-U.-A. Merits of serving UAVs via terrestrial networks: A vertical antenna radiation study. In Proceedings of the IEEE International Conference on Communications, Montreal, QC, Canada, 8–12 June 2025; pp. 1470–1475. [CrossRef]
48. Zhang, Y.; Zhao, H.; Long, Y. CMAB: A multi-attribute building dataset of China. *Sci. Data* **2025**, *12*, 430. [CrossRef] [PubMed]
49. Haklay, M.; Weber, P. OpenStreetMap: User-generated street maps. *IEEE Pervasive Comput.* **2008**, *7*, 12–18. [CrossRef]
50. Kak, A.C.; Slaney, M. *Principles of Computerized Tomographic Imaging*; IEEE Press: New York, NY, USA, 1988.
51. Otsu, N. A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man. Cybern.* **1979**, *9*, 62–66. [CrossRef]
52. AutoNavi Software Co., Ltd. Amap Open Platform. 2026. Available online: https://lbs.amap.com/ (accessed on 29 July 2026).
53. OpenStreetMap Contributors. OpenStreetMap Data. Licensed Under the Open Data Commons Open Database License (ODbL). Available online: https://www.openstreetmap.org/copyright (accessed on 5 June 2026).
54. Yapar, C.; Levie, R.; Kutyniok, G.; Caire, G. Dataset of Pathloss and ToA Radio Maps with Localization Application. In *IEEE DataPort*; IEEE: New York, NY, USA, 2022; Volume Dataset. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.