IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 11 October 2026, accepted 22 October 2026, date of publication 27 October 2026, date of current version 3 November 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3748055

RESEARCH ARTICLE

Electromagnetically Consistent Evaluation of Federated RIS Channel Estimators: Hidden Idealizations and a Full-Wave Virtual Testbed

NURI GÖKMEN KARAKIRAZ[1], AGAH OKTAY ERTAY[2], (Member, IEEE), UMUT ŞÜKRÜ YAŞAR[3], HALIT NAZLI[1], AND SAMAR A. S. AL-HASASNEH[4]

[1]Department of Electrical and Electronics Engineering, Istanbul Aydin University, 34295 Istanbul, Türkiye
[2]Department of Electrical and Electronics Engineering, Faculty of Engineering and Architecture, Erzincan Binali Yıldırım University, 24100 Erzincan, Türkiye
[3]Department of Motor Vehicles and Transportation Technologies, Kemah Vocational School, Erzincan Binali Yıldırım University, 24400 Erzincan, Türkiye
[4]General Education Unit, Al Wasl University, Dubai, United Arab Emirates

Corresponding author: Nuri Gökmen Karakiraz (nurigokmenkarakiraz@aydin.edu.tr)

ABSTRACT Learned reconfigurable intelligent surface (RIS) channel estimators are typically evaluated on an idealized diagonal phase-shift surface, with noise added to the channel and observations one nominal slot apart. This paper removes these idealizations and tests which conclusions persist. An electromagnetically consistent generator (a varactor-loaded cell calibrated to full-wave simulation, multiple-scattering mutual coupling, spherical wavefronts) is driven by an $M+1$ pilot sweep under non-Markovian Jakes dynamics; a federated long short-term memory estimator is compared with classical filters, including a deployable Kalman filter. Two idealizations systematically favor the learned estimator. Referring the cell reflection to the wave impedance instead of the element self-impedance models the passive surface as active and inflates the least-squares error 9- to 87-fold across five cell families under a dipole kernel, against about threefold for the learned estimator. Spacing observations by one slot overstates temporal correlation; retrained for the corrected interval, the recurrent estimator has 57% higher NMSE than its nominal-slot counterpart. Under the corrected models the deployable Kalman filter leads except at small normalized Doppler. Despite a 3.6- to 6.6-fold column-NMSE lead over least squares, the learned estimator delivers less field at the user in all 60 combinations of RIS configuration, beamforming or RIS-control strategy and training seed: its residual survives the matched-filter projection and is mostly amplitude error, so error-phase coherence predicts the delivered field better than NMSE. A virtual over-the-air testbed, which adds hardware-impairment models to an 84-port full-wave network with multipath, confirms the passivity audit and the ordering of simulation-trained estimators.

INDEX TERMS Channel estimation, electromagnetic modeling, federated learning, full-wave simulation, Kalman filter, multiport networks, mutual coupling, passivity, reconfigurable intelligent surface, 6G.

I. INTRODUCTION

Reconfigurable intelligent surfaces (RIS) reshape the propagation environment by applying controllable phase shifts through a large number of nearly passive elements [1], [2], [3], extending coverage at low hardware cost [2], [3] and improving energy efficiency [4]. The gains depend on channel state information (CSI): the base station (BS) must know the cascaded BS–RIS–user channel, which cannot be sensed directly because the elements carry no RF chains. Estimating it needs a training sweep over reflection patterns, so the estimation dimension and the pilot overhead both grow with the element count M [5], [6].

Least squares (LS) and linear minimum mean square error (LMMSE) estimators either degrade at low signal-to-noise ratio (SNR) or require second-order channel and noise

The associate editor coordinating the review of this manuscript and approving it for publication was Yi Ren.

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License.
For more information, see https://creativecommons.org/licenses/by/4.0/

statistics [7]. LS ignores the channel statistics, so its error grows as the SNR falls. LMMSE is the MSE-optimal linear estimator only when the channel covariance is known, and inverting it costs $\mathcal{O}(N^3(M+1)^3)$ operations [7]; in practice the covariance must be estimated, and a full-rank sample estimate needs at least $N(M+1)$ observations of the cascaded channel; in our setting even a 16-dimensional channel needs 128 observations before LMMSE outperforms LS (Section VI-A). Learned estimators address these limitations: convolutional networks denoise LS estimates [7], neural receivers estimate and detect jointly [8], and CNN-LSTM architectures use spatial and temporal structure together [9]. All of them are trained centrally, so training them on field observations requires aggregating those observations at a server. Channel observations are tied to user location and mobility, so this aggregation costs uplink capacity at scale [10] and raises privacy concerns [11].

Federated learning (FL) [12] keeps data where it is generated and shares only model updates. It has been applied to channel estimation for massive MIMO and RIS-assisted systems with a CNN [10]; multiple RISs have been used to improve over-the-air FL model aggregation [13]; and FL for physical-layer design has been reviewed more broadly [14]. Four issues remain under-addressed. First, the federated channel estimators reviewed in Section II are snapshot CNNs [10] or recurrent models evaluated without a classical sequence baseline [30], while the classical sequence filters that exploit temporal correlation (Kalman-type recursions) are the baseline against which any temporal claim must be tested. Second, FL performance is sensitive to hyperparameters, manual tuning is at odds with the automation distributed deployments require, and automated tuning fails when the tuning signal does not generalize to the full federated schedule. Third, evaluations idealize the data pipeline: noise added directly to the effective channel, IID client partitions, genie labels left implicit. Fourth, the RIS itself is abstracted as an ideal diagonal phase-shift matrix $\boldsymbol{\Theta} = \mathrm{diag}(e^{j\theta_1}, \ldots, e^{j\theta_M})$, whereas a physical metasurface has unit-cell amplitude–phase coupling and ohmic loss [15], [16], [17], an incidence-dependent response [16], mutual coupling that intensifies as the element spacing shrinks below half a wavelength [15], [16], [17], [18], and frequency dispersion [19]; a varactor-loaded cell also has a limited phase range and a polarization-dependent response (Section III-B). Electromagnetically compliant multiport and surface-impedance models [15], [16] and physics-based end-to-end channel models [19] exist, but what they imply for channel estimation, and for learned federated estimators in particular, is rarely quantified.

This paper addresses these four issues and reports how a learned federated estimator performs once the idealizations are removed. That estimator, FL-LSTM, is a recurrent model trained by federated averaging, with hyperparameters settled by an origin-blind playoff between a cost-aware Pareto search, an equal-budget random search, and a hand-tuned reference. The playoff selects the hand-tuned reference. FL-LSTM is an instance of an existing estimator class, not a new method.

Under the corrected models it is not the most accurate estimator. A deployable Kalman filter (observation-parameterized rather than oracle-tuned; "deployable" refers to the simulated receiver model) leads it at every pilot SNR, and least squares overtakes it above 5 dB. That result is itself the finding of interest, not an incidental shortcoming to be worked around. The estimator's role is instrumental: an idealization that favors learning changes the ranking only when a learned estimator is in the loop, and the advantage that appears under the idealization and vanishes without it is the quantity we report. This paper therefore does not offer a better RIS channel estimator; it quantifies how strongly idealized evaluation assumptions can create an apparent learning advantage in the setting studied here.

The contributions are the following.

1) *A pilot-faithful, electromagnetically consistent evaluation framework.* Each surface element is a varactor-loaded cell, characterized and calibrated in HFSS, coupled through a multiple-scattering model whose coupling operator is that of the multiport model of [15] and is passive by construction, with spherical wavefronts. Every observation comes from an explicit $M+1$-pattern pilot sweep with cascaded LS, under non-Markovian Jakes dynamics and non-IID federated sites (Sections III and V).
2) *Two modeling idealizations that favor learned estimators.* Referring the cell reflection to the wave impedance instead of the element self-impedance makes the passive surface active and inflates the LS error far more than the learned one, by 9 to 87× against 2.3 to 3.1× across five cell families under the dipole kernel. Spacing observations by a nominal slot instead of a full sweep overstates temporal correlation; retrained for the corrected interval, the recurrent estimator has 57% higher NMSE than its nominal-slot counterpart, against 5% for a deployable Kalman filter. Neither produces a numerical symptom (Sections VII-B, VII-H and VI-D).
3) *A quantified mechanism for why channel NMSE misranks estimators.* A learned estimator that leads LS 8.9- to 10.4-fold in per-column NMSE at half-wavelength spacing delivers less field at the user (Section VII-D), and a 3.6- to 6.6-fold lead does the same in all 60 combinations of four RIS configurations, five beamforming and RIS-control strategies and three training seeds. Its residual survives the matched-filter projection that configuration applies and is largely amplitude error that a phase-only surface ignores; the error-phase coherence predicts the delivered field, and the same structure appears at link level under MRT, ZF and RZF (Sections VII-I and VIII-B).

4) *Full-wave evidence from component to scene.* Beyond the calibrated cell and a coupling kernel extracted from the same patch geometry, a virtual over-the-air testbed solves a 32-element surface, a base-station array, moving-user probes and a blocker as one 84-port network, on which the model and the estimator ordering are re-tested with hardware impairments and multipath (Section VII-J).

The federated LSTM is not among the contributions: it is an existing estimator class and the instrument that detects the idealizations, its hyperparameters chosen by an origin-blind playoff in which a hand-tuned configuration outperforms a cost-aware Pareto search (Section VI-B). Deployable and oracle classical filters make the negative results interpretable.

Section II reviews related work, Section III the system, channel and pilot models, Section IV the federated learning and selection framework, and Section V the experimental protocol. Results are reported in four parts: estimator-level results (Section VI), electromagnetic and physical validation (Section VII), link-level and deployment implications (Section VIII), and sensitivity analyses and limitations (Section IX).

The scope is deliberately narrow. We do not establish a cryptographically secure federated protocol or claim superiority across all RIS channel models. Data locality is meant literally: raw observations stay at the site that acquired them and only model parameters reach the aggregation server. Formal protection against leakage from those updates needs differential privacy [20] or secure aggregation and lies outside this study.

Notation: Matrices are denoted by bold uppercase letters ($\mathbf{A}$), column vectors by bold lowercase letters ($\mathbf{a}$), and scalars by italic letters (a). $(\cdot)^T$ and $(\cdot)^H$ are the transpose and Hermitian transpose, $\mathrm{diag}(\cdot)$ forms a diagonal matrix, $\|\cdot\|$ the Euclidean norm, $\odot$ the Hadamard product, $J_0(\cdot)$ the zeroth-order Bessel function of the first kind, and $\mathcal{CN}(0,\sigma^2)$ a circularly symmetric complex Gaussian distribution.

II. RELATED WORK

A. CHANNEL ESTIMATION FOR RIS-AIDED SYSTEMS

Early pilot-based RIS channel estimators paired LS [5] or LMMSE [6] formulations with designed reflection patterns. Kundu and McKay [7] showed that CNN denoising of LS estimates, designed to approximate the MMSE estimator, outperforms both LS and LMMSE without explicit knowledge of the channel covariance. Deep learning has further been used for channel estimation in massive MIMO [21] and jointly with detection in OFDM receivers [8]; for large surfaces, compressive-sensing and deep-learning solutions have been proposed that configure the reflection from channels sensed at a few active elements [22]. Closest to our setting, Nguyen et al. [9] combined CNN feature extraction with LSTM temporal modeling for RIS-NOMA systems and demonstrated gains over pure-CNN and pure-LSTM estimators; their model, however, is trained centrally on pooled user data. On the classical side, when the channel evolves in time, sequence-aware filters (Kalman-type recursions built on a first-order state model) are the natural non-learned competitor; evaluations that omit them risk crediting deep models for gains that classical filtering already provides. We include both oracle-parameterized and self-parameterizing (deployable) Kalman baselines. Two technical conclusions follow for the design of a fair comparison. First, the learned estimators of [7], [9] exploit structure that the linear estimators do not use: the non-Gaussian channel distribution beyond the second-order statistics used by LMMSE in [7], which reports gains over LS and LMMSE, and the temporal correlation in [9], whose model is compared only with other learned models. A temporal learned estimator therefore has to be compared with a classical estimator that exploits the same temporal correlation; none of the learned-estimation works in Table 1 does so. Second, all of them treat the surface as an ideal phase shifter, and the pilot sweep is either not costed or counted as a rate overhead [7], never charged to the interval between observations of a time-varying channel.

B. ELECTROMAGNETICALLY CONSISTENT RIS MODELS

The communication-theoretic RIS abstraction, an ideal diagonal phase-shift matrix, discards the electromagnetics of the surface. Gradoni and Di Renzo [15] introduced an end-to-end mutual-coupling-aware model in which the RIS is a multiport network of loaded thin-wire scatterers; for single-antenna terminals in the far field the channel is proportional to $\mathbf{Z}_{\mathrm{RT}} - \mathbf{Z}_{\mathrm{RS}}(\mathbf{Z}_{\mathrm{RIS}} + \mathbf{Z}_{\mathrm{SS}})^{-1}\mathbf{Z}_{\mathrm{ST}}$, where $\mathbf{Z}_{\mathrm{SS}}$ is the mutual-impedance matrix of the scatterers and $\mathbf{Z}_{\mathrm{RIS}}$ the diagonal matrix of their tunable loads. In the rest of this paper these two matrices are written $\mathbf{Z}_{\mathrm{RIS}}$ and $\mathbf{Z}_{\mathrm{L}}$, respectively. Di Renzo et al. [16] review this model alongside surface-electromagnetic models (impedance boundary conditions, Green's functions) and their use in network optimization, and Di Renzo and del Hougne [23] survey the application of multiport network theory to modeling and optimizing reconfigurable metasurfaces. Qian and Di Renzo [17] showed that optimization which ignores mutual coupling loses substantial received power when the elements are spaced below half a wavelength, and Zheng et al. [18] demonstrated, through a misspecified Cramér–Rao bound analysis, that neglecting mutual coupling degrades RIS channel estimation, increasingly so for dense element layouts. Coupling-aware *estimation* has since become an active line in its own right: Zheng et al. [24] fold the coupling into a compressed-sensing formulation for active-RIS MIMO links and reduce its cost with a two-stage scheme in which a coarse coupling-unaware estimate drives a dictionary reduction, Qiu et al. [25] extend coupling-aware estimation to multi-user mmWave MISO systems with a three-stage uplink protocol, and Dkhan et al. [26] build mutual coupling into near-field reduced-subspace least-squares estimation through the spatial correlation matrices that define the subspace. Not all of this literature treats coupling as a penalty: Nerini et al. [27] derive globally optimal

closed-form solutions for beyond-diagonal (fully- and tree-connected) surfaces under coupling and prove that, when the optimization accounts for it, coupling raises the average channel gain under independent Rayleigh fading; they note that this result gives no insight into conventional diagonal surfaces such as the one modeled here. It is complementary to the degradation we quantify, since the two ask different questions, one about the optimized surface and one about the estimator that has to configure it. PhysFad [19] provides a physics-based end-to-end channel model built on a first-principles coupled-dipole formalism, in which each polarizable dipole is driven by the external field plus the fields radiated by all other dipoles, the same multiple-scattering structure as a Foldy–Lax formulation. These works establish the models. What remains open, and is addressed here, is what they imply for temporal, federated, learned channel estimation: how accuracy degrades with element spacing under coupling, what an idealized-model-trained network loses under realistic electromagnetics, and whether NMSE gains survive translation to surface currents, radiation patterns and field focusing. Three technical conclusions carry over. The multiport form of [15] fixes which impedance the reflection inside the coupling inverse must be referred to, the element self-impedance rather than the wave impedance, which is the basis of the audit in Section VII-B. The coupling-aware estimators [24], [25], [26] assume the coupling is known and estimate snapshots centrally with classical algorithms, so they do not say whether coupling reorders learned and classical temporal estimators. Finally, [18] shows that ignoring coupling degrades estimation increasingly with density, which motivates quantifying that degradation per estimator rather than for a single estimator.

C. FEDERATED LEARNING FOR THE PHYSICAL LAYER

FedAvg [12], together with work on its open challenges [11], its communication efficiency [28] and its degradation under non-IID data [29], established FL as a communication-efficient alternative to centralized training that keeps raw data local. Elbir and Coleri [10] introduced federated CNN training for channel estimation in conventional and RIS-assisted massive MIMO; multi-RIS-aided over-the-air FL has also been studied [13], and FL for physical-layer design has been surveyed [14]. The channel estimator of [10], which is also the one reviewed in [14], is a snapshot-based CNN, and [13] assumes that CSI is available at the base station; client heterogeneity is typically limited. Closest to the present study, Li et al. [30] treat time-varying RIS-assisted over-the-air FL with joint channel estimation and dynamics-aware user grouping, using gated recurrent units to track temporal channel structure under imperfect CSI and user heterogeneity. That work shares the recurrent-temporal and RIS-assisted-FL ingredients with ours, including federated training of the estimator on local pilot observations, and differs in how aggregation is realized and in the fidelity of the surface model. There, the shared backbone of a personalized GRU-based estimator is aggregated over the air (the user-specific heads stay local), and its estimates in turn drive that aggregation and the configuration of an ideal phase-shift surface. Here, a single global estimator is trained by digital, sample-size-weighted averaging on strictly site-local pilot observations, acquisition is pilot-faithful over an explicitly budgeted observation interval, and the surface is represented by an electromagnetically consistent multiport model with mutual coupling, cross-checked against a full-wave unit cell and validated at the level of surface currents and radiated fields. With that qualification, the temporal dimension of fading channels has still received limited attention in *data-local* federated RIS channel estimation, and we are not aware of a treatment that combines it with an electromagnetically consistent surface model and pilot-faithful acquisition; the comparison against classical sequence filters is also typically absent. The technical conclusion is that the federated estimators to date are either snapshot CNNs [10] or recurrent models evaluated on an ideal surface without a classical sequence baseline [30], so neither shows whether federated temporal learning outperforms a classical filter on the same observations.

TABLE 1. Features of closely related work.

Ref.	C1	C2	C3	C4	C5	C6	C7	C8
[7]	–	–	✓	–	–	–	✓	–
[9]	–	–	–	○	–	–	–	–
[10]	–	–	○	–	✓	–	–	–
[18]	○	✓	○	–	–	–	–	–
[22]	–	–	○	–	–	–	✓	–
[24]	–	✓	○	–	–	–	✓	–
[25]	–	✓	○	–	–	–	–	–
[26]	–	✓	○	–	–	–	–	–
[30]	–	–	○	✓	✓	–	✓	–
This work	✓	✓	✓	✓	✓	✓	✓	✓

C1 element response beyond an ideal phase shift; C2 mutual coupling; C3 pilot cost charged to the coherence time or the rate; C4 time-varying channel; C5 data-local (federated) training; C6 classical sequence-filter baseline; C7 evaluation at configuration, field or link level; C8 full-wave or measured validation. ✓: yes; ○: partly (C3: pilots counted but not charged); –: no.

D. HYPERPARAMETER OPTIMIZATION

PSO [31], [32] is a population-based optimizer that can also be applied to neural hyperparameter search, alongside random search [33] and bandit-based methods such as Hyperband [34], which formalize the trade-off between evaluation fidelity and search breadth. In federated settings the fidelity problem is acute: a configuration that performs well after a few rounds may underperform over the full schedule. Our design differs from a generic evolutionary search in evaluating fitness with federated rather than centralized proxy runs, in optimizing communication payload and local computation as explicit objectives alongside validation NMSE, in reporting the search's own run-to-run variability against an equal-budget random-search competitor, and in resolving the final selection with a high-fidelity playoff that includes hand-tuned and warm-start references.

The technical conclusion drawn from [34] is that low-cost proxy evaluations trade fidelity for breadth, so a search over federated hyperparameters has to verify that its proxy ranks candidates as the full schedule does; Section VI-B tests exactly that.

What is new. Table 1 shows that the individual ingredients exist; none of the nine works satisfies more than three of the eight features in full, and none compares a learned estimator with a classical sequence filter. We do not claim the ingredients, and the federated LSTM is an existing estimator class. The methodological contribution lies in three capabilities that none of the ingredients provides alone. The first is a pair of consistency audits, the impedance reference of the coupling inverse and the observation interval of a swept pilot, each tied to a quantified change in the relative estimator ranking. Either error can be found without a learned estimator, by a passivity check or from the pilot timing, but its consequence for learning-based evaluation shows only when learned and classical estimators are compared under the same corrected model. The second is an evaluation chain from channel estimate to bias commands, coupled currents, delivered field and link, together with a quantified mechanism (error alignment with the matched filter, amplitude and phase split, error-phase coherence) that predicts when channel NMSE misranks estimators. The third is full-wave grounding at component and system level: a calibrated unit cell, a kernel extracted from the same geometry, and a virtual over-the-air testbed on which the ranking is re-tested with hardware impairments.

Each of these four directions has been studied on its own. The present work sits at their intersection, under an evaluation protocol designed to address the standard criticism of each.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. NETWORK GEOMETRY AND CHANNELS

We consider a multi-user MISO system in which a BS with an N-element uniform linear array (ULA) serves K single-antenna users with the aid of an RIS comprising M passive reflecting elements (Fig. 1). The system operates in TDD; uplink pilots yield the CSI used for downlink precoding by reciprocity. In the uplink, the effective SIMO channel of a user aggregates the direct and RIS-cascaded links:

$$\mathbf{h}_{\text{eff}} = \mathbf{h}_{\text{d}} + \mathbf{H}_{\text{br}}\,\text{diag}(\boldsymbol{\theta})\,\mathbf{h}_{\text{r}} \in \mathbb{C}^{N}, \tag{1}$$

where $\mathbf{h}_{\text{d}} \in \mathbb{C}^{N}$ is the direct user–BS channel, $\mathbf{H}_{\text{br}} \in \mathbb{C}^{N\times M}$ the RIS–BS channel, $\mathbf{h}_{\text{r}} \in \mathbb{C}^{M}$ the user–RIS channel, and $\boldsymbol{\theta} = [e^{j\theta_1}, \ldots, e^{j\theta_M}]^{T}$, $\theta_m \in [0, 2\pi)$, the RIS phase configuration. The downlink channel is the transpose $\mathbf{h}_{\text{eff}}^{T}$ by TDD reciprocity; downlink transmission with linear precoding $\mathbf{x} = \mathbf{W}\mathbf{s}$ gives user k

$$y_k = \mathbf{h}_{\text{eff},k}^{T}\mathbf{W}\mathbf{s} + n_k, \qquad n_k \sim \mathcal{CN}(0, \sigma^2). \tag{2}$$

All links carry distance-dependent path loss with exponents α_{d} (direct, obstructed), α_{br} (BS–RIS, LoS-dominated), and α_{ru} (RIS–user). The RIS-related links are Rician with steering-vector LoS components,

$$\mathbf{H} = \sqrt{\beta}\left(\sqrt{\tfrac{K_R}{K_R+1}}\,\mathbf{H}^{\text{LoS}} + \sqrt{\tfrac{1}{K_R+1}}\,\mathbf{H}^{\text{NLoS}}\right), \tag{3}$$

where β is the path gain, $\mathbf{H}^{\text{LoS}}$ is built from ULA steering vectors of the actual geometry (e.g., $\mathbf{H}_{\text{br}}^{\text{LoS}} = \mathbf{a}_N(\phi_{\text{br}})\mathbf{a}_M(\phi_{\text{rb}})^{H}$), and $\mathbf{H}^{\text{NLoS}}$ has IID $\mathcal{CN}(0, 1)$ entries. The direct link is Rayleigh with path loss. The BS–RIS channel $\mathbf{H}_{\text{br}}$ is static (both terminals fixed); the user-side links evolve in time as described in Section III-C. Equation (1) with unit-modulus $\boldsymbol{\theta}$ is the idealized RIS abstraction; the electromagnetically consistent response that replaces it is developed next.

B. ELECTROMAGNETICALLY CONSISTENT RIS MODEL

A physical RIS element is not an ideal phase shifter. Three layers of electromagnetics are added to the channel generator, each optional so that the idealized model remains available as a controlled reference.

1) UNIT-CELL SURFACE IMPEDANCE

Each element is a grounded, varactor-loaded cell described by a tunable surface impedance formed by the patch inductance L_p in parallel with the series varactor branch,

$$Z_s(f, u) = \left[\frac{1}{j\omega L_p} + \frac{1}{R_v + \frac{1}{j\omega C_v(u)}}\right]^{-1},$$
$$C_v(u) = C_{\min} + (C_{\max} - C_{\min})\,u, \tag{4}$$

where $u \in [0, 1]$ is the normalized bias state and R_v aggregates varactor and conductor loss. Under the impedance boundary condition, oblique incidence at angle θ_i gives the polarization-dependent reflection coefficient

$$\Gamma_{\text{TE}} = \frac{Z_s - \eta_0/\cos\theta_i}{Z_s + \eta_0/\cos\theta_i}, \qquad \Gamma_{\text{TM}} = \frac{Z_s - \eta_0\cos\theta_i}{Z_s + \eta_0\cos\theta_i}. \tag{5}$$

Sweeping the bias through the resonance of (4) traces the classical reflectarray S-curve: at $f_c = 3.5$ GHz the achievable phase range is ≈320° (not 360°), with a loss dip of ≈3.3 dB at resonance, so amplitude and phase are not independently controllable, as expected for a lossy resonant varactor-loaded cell. The response is dispersive in f and depends on incidence angle and polarization; all of these are characterized in Section VII-A. The controller holds a lookup table $u^{\star}(\theta)$ calibrated at normal incidence: a commanded phase θ_m is mapped to the bias $u^{\star}(\theta_m)$ (commands in the unreachable sector clip to the nearest achievable endpoint), and the realized element *scattering response* is $\Phi_m = A_m e^{j\phi_m} = \Gamma_{\text{TE}}(f_c, 0, u^{\star}(\theta_m))$.

2) MUTUAL COUPLING AS MULTIPLE SCATTERING

At element spacings $d \lesssim \lambda/2$ the elements interact: the field exciting element m includes the fields re-scattered by all other elements. Two distinct reflection quantities enter this description, and they must not be conflated because they are referenced to *different* impedances. The scattering response Φ_m of (5) is referenced to the wave impedance $\eta_0/\cos\theta_i \approx 377\,\Omega$ and sets how strongly element m re-radiates the field

incident upon it. The coupling kernel, by contrast, is normalized to the element self-impedance, so the reflection that enters the multiple-scattering inverse must carry the same reference:

$$\varrho_m = \frac{Z_{11}}{Z_{11} + Z_s\big(f, u^\star(\theta_m)\big)}. \tag{6}$$

With $\mathbf{\Phi} = \mathrm{diag}(\Phi_1, \ldots, \Phi_M)$, $\boldsymbol{\varrho} = \mathrm{diag}(\varrho_1, \ldots, \varrho_M)$ and a coupling kernel $\mathbf{K}$, the self-consistent (Foldy–Lax) solution replaces the diagonal response by

$$\mathbf{\Phi}_{\mathrm{eff}} = \big(\mathbf{I} + \boldsymbol{\varrho}\mathbf{K}\big)^{-1}\mathbf{\Phi} \in \mathbb{C}^{M\times M}, \tag{7}$$

which is full (not diagonal): each incident excitation couples into all elements. In the uncoupled limit $\mathbf{K} = \mathbf{0}$, (7) returns $\mathbf{\Phi}$ exactly. The kernel entries are the induced-EMF mutual-to-self impedance ratios of the canonical thin resonant dipole, $K_{mn} = Z_{mn}(d_{mn})/Z_{11}$ ($K_{mm} = 0$), computed from the classical closed forms (cosine/sine-integral expressions; $Z_{11} = 73.1 + j42.5\,\Omega$, $Z_{21}(\lambda/2) = -12.5 - j29.9\,\Omega$). The coupling operator of (7) is that of the electromagnetically compliant multiport impedance model of [15] (in our notation; the mutual-impedance and load matrices are $\mathbf{Z}_{\mathrm{SS}}$ and $\mathbf{Z}_{\mathrm{RIS}}$ in [15]): with $\mathbf{Z}_{\mathrm{RIS}} = Z_{11}(\mathbf{I} + \mathbf{K})$ and the cell impedances acting as the tunable loads $\mathbf{Z}_{\mathrm{L}} = \mathrm{diag}(Z_{s,1}, \ldots, Z_{s,M})$, the factorization $\mathbf{Z}_{\mathrm{RIS}} + \mathbf{Z}_{\mathrm{L}} = \mathrm{diag}(Z_{11} + Z_{s,m})\,(\mathbf{I} + \boldsymbol{\varrho}\mathbf{K})$ shows that the element currents

$$\mathbf{i} = \big(\mathbf{Z}_{\mathrm{RIS}} + \mathbf{Z}_{\mathrm{L}}\big)^{-1}\mathbf{v}_{\mathrm{inc}}, \tag{8}$$

are exactly the uncoupled currents left-multiplied by $(\mathbf{I} + \boldsymbol{\varrho}\mathbf{K})^{-1}$; this identity is verified numerically to machine precision. Formulation (7) applies the same operator to the calibrated cell response $\mathbf{\Phi}$, which takes the place of the loaded-dipole response $\mathrm{diag}\big((Z_{11} + Z_{s,m})^{-1}\big)$ of [15]. Writing (7) with $\mathbf{\Phi}$ in place of $\boldsymbol{\varrho}$ inside the inverse, that is, substituting a 377 Ω-referenced quantity into a Z_{11}-referenced expression, is not a benign approximation: inverting the identification then implies element loads $Z_{\mathrm{L},m} = Z_{11}(1/\Phi_m - 1)$ with *negative* resistance over 91–95% of the bias range (as low as $-170\,\Omega$), i.e. an active surface. In the present cell that misreference drives $(\mathbf{I} + \mathbf{\Phi}\mathbf{K})$ toward singularity and inflates $\|\mathbf{\Phi}_{\mathrm{eff}}\|_2$ to as much as 893. Under (6) the problem cannot arise: the cell is a passive RLC network ($\Re Z_s \geq 0$) and the induced-EMF radiation-resistance matrix is positive semidefinite across the swept spacings, so $\mathbf{Z}_{\mathrm{RIS}} + \mathbf{Z}_{\mathrm{L}}$ is positive-real and no element can supply power. Passivity holds by construction rather than by numerical accident. The closed-form kernel itself is validated in two independent ways in Section VII-A: against direct numerical quadrature of the sinusoidal-current near-field EMF integral, and against full-wave two-port HFSS simulations of loaded dipole pairs across d/λ.

3) NEAR-FIELD (SPHERICAL-WAVE) PROPAGATION

For an aperture of M elements the far-field distance $2D^2/\lambda$ can exceed the user distances of interest, so plane-wave steering vectors are replaced by per-element spherical propagation. With element positions $\mathbf{r}_m$ and BS antenna positions $\mathbf{b}_n$, the deterministic LoS factors are

$$[\mathbf{G}_{\mathrm{br}}]_{nm} \propto \frac{e^{jk_0\|\mathbf{b}_n - \mathbf{r}_m\|}}{\|\mathbf{b}_n - \mathbf{r}_m\|}, \qquad [\mathbf{g}_{\mathrm{ru}}]_m \propto \frac{e^{jk_0\|\mathbf{r}_m - \mathbf{r}_{\mathrm{UE}_k}\|}}{\|\mathbf{r}_m - \mathbf{r}_{\mathrm{UE}_k}\|}, \tag{9}$$

normalized to unit mean power so that the large-scale path-loss model of (3) is preserved (a linearized far-field variant is retained as a switch for comparison). The electromagnetically consistent effective channel replacing (1) is then

$$\mathbf{h}_{\mathrm{eff}}^{\mathrm{EM}} = \mathbf{h}_{\mathrm{d}} + \mathbf{G}_{\mathrm{br}}\,\mathbf{\Phi}_{\mathrm{eff}}\,\mathbf{g}_{\mathrm{ru}}, \tag{10}$$

with the Rician NLoS and Jakes temporal structure of Sections III-A and III-C applied unchanged to the stochastic components.

4) MODEL HIERARCHY AND ESTIMATOR-SIDE MISMATCH

Three nested models are used throughout: *ideal* ($\mathbf{\Phi}_{\mathrm{eff}} = \mathrm{diag}(e^{j\theta_m})$, the conventional abstraction), *EM-nc* (realized unit-cell response (5), no coupling), and *EM-mc* (full model (7)). Because the three share all random draws (geometry, fading, noise), differences isolate the RIS physics. The pilot protocol of Section III-D is left unchanged: the LS estimator inverts the ideal linear cascade (14) while the data are generated from the nonlinear coupled response of (10), the systematic model error a deployed estimator would face. Whether that error is of a kind classical filters must inherit and label-driven learning can correct is an empirical question, and Sections VII-A–VII-D answer it: under a passivity-consistent coupling model the answer is largely no, in contrast to what the same experiments report when the coupling model is misreferenced.

C. TEMPORAL DYNAMICS: JAKES SUM-OF-SINUSOIDS

User motion makes $\mathbf{h}_{\mathrm{d}}$ and $\mathbf{h}_{\mathrm{r}}$ time-varying over the T observations of a window. The interval between two consecutive observations is fixed by the pilot protocol, not by the nominal frame structure: one observation costs a complete sweep of $P = M+1$ reflection patterns (Section III-D), so

$$T_{\mathrm{obs}} = P\,T_p = (M+1)\,\frac{1 + \gamma_{\mathrm{cp}}}{\Delta f}. \tag{11}$$

At $\Delta f = 15\,\mathrm{kHz}$ with a normal cyclic prefix, $T_p = 71.35\,\mu\mathrm{s}$, so the $M = 32$ surface studied here separates consecutive observations by 2.355 ms, not by the 1 ms nominal slot. Substituting the slot for (11) would fit two and a third sweeps into the time of one, and at 30 km/h it would raise the lag-1 correlation from 0.546 to 0.909; every sequence estimator would then be scored on a channel that varies more slowly than the protocol allows. All temporal quantities below use T_{obs}. Note the direction of the remaining error: (11) assumes sweeps run back to back, with no data phase between them, so it is a *lower* bound on the interval a deployed system would see. If data are also transmitted between acquisitions, the channel decorrelates further, and the penalties Section VI-A attributes to the correction are correspondingly conservative.

Rather than imposing a first-order autoregressive evolution, which would make a Kalman filter structurally matched to the true dynamics, the NLoS components follow the Clarke/Jakes model, realized as a sum of sinusoids (SOS): each scalar NLoS coefficient evolves as

$$s^{(t)} = \frac{1}{\sqrt{N_p}} \sum_{p=1}^{N_p} \exp\big(j(2\pi f_D T_{\text{obs}} \cos(\alpha_p)\, t + \varphi_p)\big), \qquad (12)$$

where the sum runs over N_p paths with arrival angles α_p and phases φ_p drawn IID from $U(0, 2\pi)$, at maximum Doppler $f_D = v f_c / c$. The resulting autocorrelation is $\mathbb{E}[s^{(t+\tau)} s^{(t)*}] = J_0(2\pi f_D T_{\text{obs}} \tau)$, a non-Markovian process for which no first-order recursion is exact. The LoS component rotates in phase at the LoS Doppler shift, $\mathbf{h}_{\text{r}}^{\text{LoS},(t)} = \mathbf{h}_{\text{r}}^{\text{LoS}} e^{j2\pi f_D \cos(\vartheta_L) T_{\text{obs}} t}$. The true lag-1 correlation of the process, $\rho = J_0(2\pi f_D T_{\text{obs}})$, is recorded with each sample and supplied to the oracle Kalman baseline; the first-order model

$$\mathbf{h}^{(t)} = \rho\, \mathbf{h}^{(t-1)} + \sqrt{1 - \rho^2}\, \boldsymbol{\varepsilon}^{(t)} \qquad (13)$$

is thus *only* the approximation used by the Kalman baselines, exactly the model mismatch a deployed first-order tracker would face. (For reproducibility the codebase retains a pure-AR(1) mode in which (13) is the true dynamics.) In the temporal-correlation ablation, a target lag-1 value ρ is converted to an equivalent Doppler through the inverse of J_0, so both dynamic models are swept on the same physical axis, and each ρ maps to an equivalent user speed. That conversion also uses (11), so a quoted equivalent speed is specific to the surface size at which it was quoted.

One consequence of (11) has to be stated before any result is read: T_{obs} grows linearly with M. The M-sweep of Section VI-F therefore does not isolate aperture size. A larger surface is harder in two ways at once, because it raises the estimation dimension *and* because its longer sweep lets the channel decorrelate further between observations, and the two effects cannot be separated within this protocol. We read that sweep as the joint cost of surface size under a full element-wise pilot protocol, which is the quantity a system designer faces, rather than as an effect of aperture alone. Holding T_{obs} fixed across M is possible only by shortening the sweep, which changes the estimator's input, or by lowering the speed with M, which changes the channel; either way the comparison would no longer be at constant protocol.

D. PILOT-BASED CASCADED CHANNEL ESTIMATION

The estimator input is not the true channel plus noise; it is the output of an explicit pilot protocol executed once per observation. Define the cascaded channel matrix

$$\mathbf{G}^{(t)} = \big[\, \mathbf{h}_{\text{d}}^{(t)},\ \mathbf{H}_{\text{br}} \operatorname{diag}(\mathbf{h}_{\text{r}}^{(t)}) \,\big] \in \mathbb{C}^{N \times (M+1)}, \qquad (14)$$

whose first column is the (unmodulated) direct path. During observation t, the RIS sweeps $P = M+1$ training configurations $\mathbf{v}_p = [1; \boldsymbol{\theta}_p]$ taken as the columns of the $P \times P$ DFT matrix $\mathbf{F}$ (first row all ones, matching the non-configurable direct path), and the BS receives $\mathbf{Y}^{(t)} = \mathbf{G}^{(t)}\mathbf{F} + \mathbf{N}^{(t)}$ with IID $\mathcal{CN}(0, \sigma_p^2)$ noise. Since $\mathbf{F}\mathbf{F}^H = P\,\mathbf{I}$, the LS estimate is

$$\widehat{\mathbf{G}}^{(t)} = \frac{1}{P}\, \mathbf{Y}^{(t)} \mathbf{F}^H, \qquad (15)$$

and the per-slot observation passed to the estimators is the effective channel synthesized under the operating phase configuration $\boldsymbol{\theta}_d$:

$$\hat{\mathbf{h}}^{(t)} = \widehat{\mathbf{G}}^{(t)}[1; \boldsymbol{\theta}_d], \qquad \hat{\mathbf{h}}^{(t)} = \mathbf{h}_{\text{eff}}^{(t)} + \mathbf{e}^{(t)}, \qquad (16)$$

with estimation-error variance $\sigma_e^2 = \sigma_p^2$ per antenna (the $1/P$ LS noise averaging is exactly consumed by the P-term recombination). The pilot noise power is set per sample from the pilot SNR, $\sigma_p^2 = (\|\mathbf{G}\|_F^2 / N) \cdot 10^{-\text{SNR}/10}$. Three consequences run through the results, all visible in the M-sweep of Section VI-F: the overhead is $P = M+1$ symbols per observation, the estimation problem grows with the surface, and the observations are spaced $T_{\text{obs}} = P\,T_p$ apart, so the channel decorrelates further between them as M grows.

1) COHERENCE BUDGET OF THE PILOT SWEEP

Counting pilot symbols does not capture their cost, so we convert the overhead into time and into rate. With subcarrier spacing Δf and cyclic-prefix fraction γ_{cp}, one OFDM symbol lasts $T_p = (1 + \gamma_{\text{cp}})/\Delta f$; with Clarke's rule the channel stays coherent for $T_c \approx \kappa / f_D$ with $\kappa = 0.423$, so the coherence interval holds

$$N_{\text{coh}} = \left\lfloor \frac{T_c}{T_p} \right\rfloor = \left\lfloor \frac{\kappa\, \Delta f}{(1 + \gamma_{\text{cp}}) f_D} \right\rfloor \qquad (17)$$

symbols. A full sweep is feasible only if $P = M+1 \le N_{\text{coh}}$, which caps the surface size at $M_{\max} = N_{\text{coh}} - 1$ and, equivalently, caps the user speed at $v_{\max} = 3.6\,\kappa c / \big((M+1) T_p f_c\big)$ km/h. The fraction of the interval consumed by pilots reduces the achievable rate, so all spectral efficiencies are additionally reported in *net* form,

$$R_{\text{net}} = \eta_{\text{net}} \sum_{k=1}^{K} \log_2\big(1 + \text{SINR}_k\big),$$
$$\eta_{\text{net}} = \max\Big(0,\ 1 - \tfrac{P}{N_{\text{coh}}}\Big). \qquad (18)$$

We use $\Delta f = 15$ kHz with a normal cyclic prefix and 14 symbols per slot, so the nominal NR slot is $14\,T_p = 0.999$ ms. That slot sets the frame grid and enters (17) only through T_p; it is *not* the spacing between channel observations, which is $T_{\text{obs}} = P\,T_p$ by (11) and exceeds one slot for every $M \ge 14$. Section VIII-A reports the resulting feasibility map. The protocol modeled here is the *full* element-wise sweep, the worst case; element grouping [5] and compressive designs that place a few active sensing elements on the surface [22] reduce P below $M+1$ and would relax the bound accordingly.

E. LABEL ACQUISITION

Supervised training requires target channels. Two modes are implemented and compared. In the *genie* mode the label is

the clean channel $\mathbf{h}_{\text{eff}}^{(T)}$; this corresponds to offline training on a synthetic or digital-twin data source and is stated explicitly as such: it is an upper training reference, not a deployment claim. In the *teacher* mode, labels are produced by a realistic acquisition procedure that a fielded system can execute: an occasional long-pilot transmission with R_t repetitions of the pattern sweep at a G_t higher pilot power, followed by LS (15). Because the label source is the weakest link in any supervised deployment argument, we do not confine it to a single ablation row: the *entire* main comparison is repeated under teacher labels, on identical test sets and with identical seeds, and both columns are reported side by side (Supplementary Section S3). All test-set metrics are always computed against the true channel, in both label modes.

F. FEDERATED SITES AND DATA LOCALITY

Federated training is performed across C *sites* (distributed BS/access points), which are distinct from the K simultaneously served users of the link-level evaluation. Each site runs the TDD protocol above for the users in its own coverage zone: pilot reception, LSTM inference, and precoding all occur at the site, and raw pilot observations never leave it; only model parameters traverse the backhaul. Sites are statistically heterogeneous (non-IID) with persistent per-site characteristics: an angular sector, a near or far distance ring, a site-specific Rician-factor range, a site-specific speed range (hence Doppler), a contiguous sub-range of the pilot-SNR grid, and an unequal share of the training data. This heterogeneity is drawn once per data seed and kept fixed, so every federated round sees the same non-IID partition.

Feature normalization and validation are computed in federated form: each site uploads its sample count and the first two raw moments of its labels (24 bytes), and the error sum and count on its local validation partition, from which the server forms exactly the pooled scale and the pooled metric (Supplementary Section S21).

G. PROBLEM FORMULATION

The goal is a mapping f_w from the observation sequence $(\hat{\mathbf{h}}^{(1)}, \ldots, \hat{\mathbf{h}}^{(T)})$ (stacked real and imaginary parts) to an estimate of the final-observation channel that minimizes the normalized mean square error,

$$\text{NMSE} = \mathbb{E}\left[\frac{\|\hat{\mathbf{h}}_{\text{eff}} - \mathbf{h}_{\text{eff}}\|^2}{\|\mathbf{h}_{\text{eff}}\|^2}\right], \tag{19}$$

subject to the *data-locality constraint* that raw observation sequences remain at the site that acquired them: only model parameters may be exchanged with the server. We also report $\text{NRMSE} = \sqrt{\text{NMSE}}$ (computed per run, then aggregated), the form in which [9] describes its NRMSE metric; the formal definition in [9] normalizes the RMSE by the range of the target channel amplitudes instead, so absolute values are not comparable across the two papers. Being a monotone transform, it does not change the ranking, so the curve is not plotted.

IV. FEDERATED LSTM AND HYPERPARAMETER-SELECTION FRAMEWORK

A. LSTM CHANNEL ESTIMATOR

Fading with Jakes dynamics carries information across observations that snapshot estimators discard. We adopt an LSTM [35] operating on the observation sequence, in the now-standard form that adds a forget gate $\mathbf{f}_t$ [36] to the original cell of [35]. The cell, whose standard equations are collected in Appendix A, takes the pilot estimate $\mathbf{x}_t$ of observation t as input, its gates decide what the cell state keeps from one observation to the next, and the estimate is read out from the final hidden state $\mathbf{s}_T$ as $\hat{\mathbf{h}}_{\text{eff}} = \mathbf{W}_{\text{out}}\mathbf{s}_T + \mathbf{b}_{\text{out}}$; the hidden state is written $\mathbf{s}_t$ rather than $\mathbf{h}_t$, because $\mathbf{h}$ denotes a channel throughout. The hidden width H, the learning rate η, and the local-epoch count E are fixed by the search-and-playoff procedure of Section IV-C, which on this problem returns the hand-tuned reference (Section VI-B).

B. WEIGHTED FEDERATED TRAINING

Each site c holds a local dataset $\mathcal{D}_c$ of n_c observation sequences, with n_c unequal across sites (Section III-F). Training proceeds in rounds: the server broadcasts the global weights, each participating site performs E local epochs of Adam updates on $\mathcal{D}_c$, and the server aggregates by *weighted* federated averaging [12], each local model weighted by its share of the samples of the sites participating in that round (a random fraction $q \le 1$ of all sites; $q = 1$ unless stated), as in (28) of Appendix A. The best-validation global model over all rounds is retained. Only weight vectors traverse the network; observation sequences never do. Optionally, uplink and/or downlink weights are uniformly quantized per tensor to b bits with min/max metadata (Section IV-D).

C. HYPERPARAMETER SELECTION: PARETO SEARCH AND ORIGIN-BLIND PLAYOFF

The triple $\boldsymbol{\vartheta} = (\log_{10}\eta, H, E)$ is selected in two stages; the design choices, the conservative warm start and the complete procedure (Algorithm 1) are given in Supplementary Section S7. In the first stage a multi-objective particle swarm [31], [32], [37], with the update (29)–(30), scores each candidate by the vector

$$\mathbf{f}(\boldsymbol{\vartheta}) = \left[\text{NMSE}_{\text{val}}(\boldsymbol{\vartheta}),\ C_{\text{tot}}(\boldsymbol{\vartheta}),\ F_{\text{loc}}(\boldsymbol{\vartheta})\right], \tag{20}$$

all minimized: the validation NMSE after a five-round federated proxy run, the total federated payload $C_{\text{tot}} = R\,C_{\text{rnd}}$ in bytes, and the total local training cost $F_{\text{loc}} = 3RE\,n\,\phi(H)$ in floating-point operations. It keeps an archive of non-dominated candidates pruned by the crowding distance of NSGA-II [38], and it is repeated three times and matched by three equal-budget random searches [33]. One operating point is read off the front by

$$J(\boldsymbol{\vartheta}) = \text{NMSE}_{\text{val}} + \lambda_C \frac{C_{\text{tot}}}{C_0} + \lambda_F \frac{F_{\text{loc}}}{F_0}, \tag{21}$$

with C_0, F_0 the costs of the hand-tuned reference (H=64, E=3) and $\lambda_C = \lambda_F = 0.02$. In the second stage, because

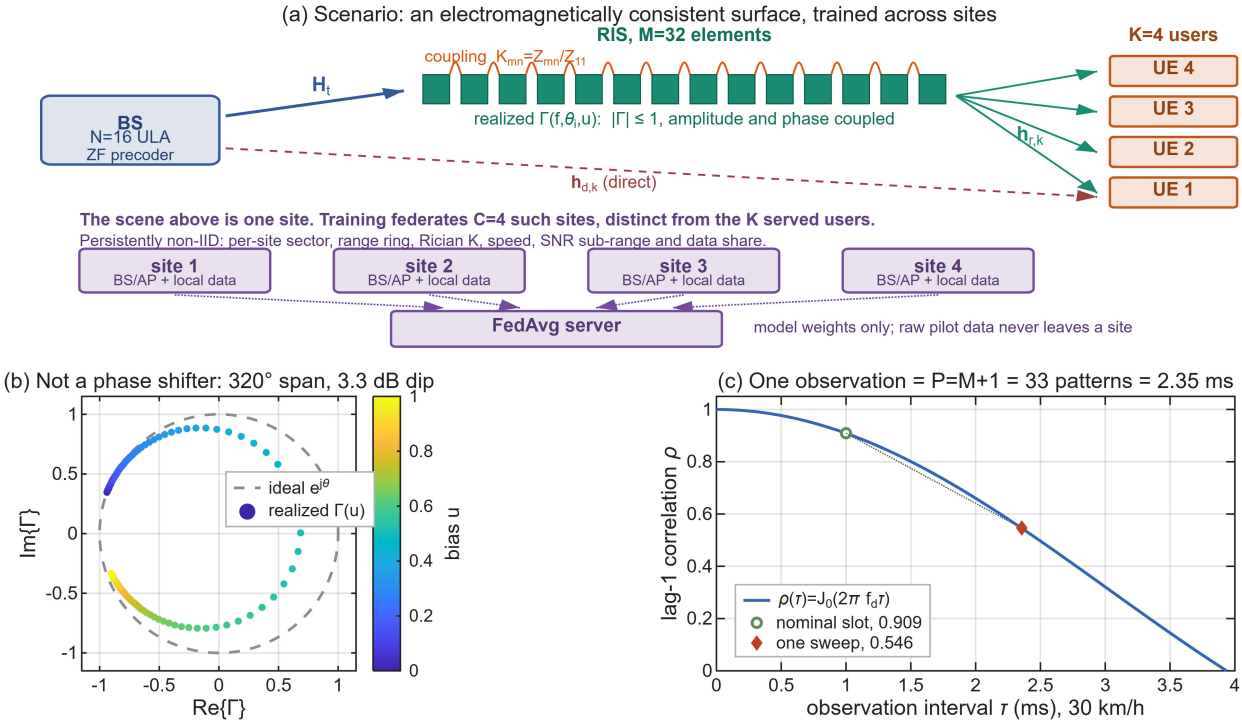


FIGURE 1. System and protocol. (a) An $N = 16$-antenna BS serves $K = 4$ users through a coupled $M = 32$-element RIS; training federates $C = 4$ sites. (b) Realized reflection locus against the ideal $e^{j\theta}$ model. (c) Time cost of one observation and its effect on temporal correlation.

intermediate losses deviate from fully trained ones [34] and low-fidelity fitness can therefore favor early progress over final accuracy, an origin-blind playoff retrains the three best distinct front points, the best random-search candidate, the hand-tuned reference ($\eta = 10^{-3}$, $H = 64$, $E = 3$) and a conservative warm start ($\eta = 10^{-3}$, $H = 16$, $E = 3$) on the full partition, and selects the minimizer of the full-fidelity objective J_{HF} regardless of which procedure proposed it.

D. MEASURED COMMUNICATION ACCOUNTING

Centralized training transports every raw sample to the server; FL transports model weights each round in both directions. Rather than estimating these costs analytically, the training loop *measures* them: every uplink and downlink transfer is counted in bytes at its actual quantization resolution, including the per-tensor min/max metadata (8 bytes per tensor) required to dequantize. Downlink cost is reported under two accountings, *unicast* (each site receives its own copy, C transmissions) and *broadcast* (one common transmission), since either may apply depending on the backhaul. For a training set of S samples of β_s bytes each, the breakeven sample count beyond which FL transfers fewer bytes than raw-data upload is $S^* = C_{\mathrm{FL}}^{\mathrm{total}}/\beta_s$, reported per scenario in Section VIII-D. Independently of byte counts, FL removes the disclosure of raw channel observations, which is the primary motivation here.

V. EXPERIMENTAL SETUP

All methods are trained on identical data generated from the model of Section III with the parameters of Table 3.

Classical references. (i) *LS*: the final-observation pilot estimate (16), the input any method receives. (ii) *Oracle scalar Wiener*: the scalar Wiener shrinkage of the final observation with *perfect* knowledge of signal and noise powers. It uses no temporal information and is *not* a lower bound for sequence estimators; it isolates the gain of perfect second-order statistics on a single snapshot. (iii) *Kalman, AR(1) oracle*: a per-antenna scalar Kalman filter with state model (13), supplied externally with the *true* lag-1 correlation $\rho = J_0(2\pi f_D T_{\mathrm{obs}})$ of each sample and the true stationary power. Under Jakes dynamics this is an upper reference (exact lag-1 correlation, mismatched first-order model) unavailable to any deployed filter. (iv) *Kalman, deployable*: the same recursion with every parameter estimated from information available at the receiver: σ_e^2 from the pilot protocol (legitimate system knowledge), signal power $\hat{\sigma}_h^2 = \max(\hat{P}_y - \sigma_e^2,\ 0.05\hat{P}_y)$ from the observations, and a *complex* transition coefficient $\hat{a} = \langle y^{(t)} y^{(t-1)*} \rangle / \hat{\sigma}_h^2$ (magnitude clipped at 0.995) that partially tracks the LoS Doppler rotation. This is the classical filter a practitioner could deploy, and hence the fair non-learned competitor.

Learned benchmarks. (v) *CNN-LSTM (centralized)*: the architecture of [9], adapted from its RIS-NOMA origin to our MU-MISO model and trained centrally on pooled data. (vi) *LSTM (centralized)*: the estimator of Section IV trained centrally with the hand-tuned configuration $(\eta, H) = (10^{-3}, 64)$, the data-locality-violating counterpart of the selected estimator. (vii) *FedCNN*: a federated CNN following [10], smaller (78 896 against 600 192 parameters) and snapshot-based as in the original design, but trained with the

same weighted federated averaging and local-epoch budget as the other federated arms, whereas [10] averages gradients uniformly at every iteration. (viii) *FL-LSTM (untuned)*: the federated pipeline with a configuration drawn uniformly at random from the PSO search space, independently per run, that is, a random-search experiment of size one in the sense of [33], which we take as the expected performance of a practitioner without tuning knowledge; Section VI-B complements this with a dedicated 15-draw risk analysis. Published numbers from [9], [10] are not transplanted into our tables, since error values (a range-normalized NRMSE in [9], NMSE in [10]) are not comparable across system models; all architectures are re-trained from scratch on identical data, and all adaptations are documented in the implementation.

Budget parity. All federated methods (selected FL-LSTM, FedCNN, untuned) use the same local-epoch count $E = E^*$, and the centralized models are trained for $R \times E^*$ epochs, equalizing data passes (with $R = 50$, $E^* = 3$: 150 centralized epochs). No method sees more data or more gradient passes than another.

Statistical protocol. Every learned result is reported over a full cross design of 10 data seeds (independent channel data *and* site geometries) $\times 3$ training seeds. These runs are *not* exchangeable: runs sharing a data seed share the test set, the geometry realization and the site distribution, so the design is nested and the effective number of independent observations is below the run count. Confidence intervals therefore come from a two-stage cluster bootstrap that resamples data seeds with replacement and then training seeds within each, reported with the intraclass correlation ICC, the design effect $1 + (n_T - 1)$ICC and the effective sample size, so the information content of the design is visible rather than implied. The naive t-interval that treats runs as independent is narrower whenever the ICC is positive, and is not used. The data seed dominates the variance, hence the 10×3 layout rather than 3×5 (also available as a reproduction mode). The main comparisons are *paired* per run: mean difference, bootstrap CI, Cohen's d, win fraction. Classical estimators are deterministic given the data seed.

Link level. Precoders are built on the estimated channels and evaluated on the true channels: zero-forcing $\mathbf{W}_{\mathrm{ZF}} = \hat{\mathbf{H}}^*(\hat{\mathbf{H}}^T\hat{\mathbf{H}}^*)^{-1}$ [39] and regularized ZF [40], whose regularization weight is the number of users divided by the data SNR, the value that maximizes the large-system SINR approximation of [40], with per-user

$$\mathrm{SINR}_k = \frac{|\mathbf{h}_k^T \mathbf{w}_k|^2}{\sum_{j \neq k} |\mathbf{h}_k^T \mathbf{w}_j|^2 + \sigma^2}, \tag{22}$$

so estimation errors surface as residual inter-user interference. The pilot SNR is held fixed at 10 dB while the data SNR is swept, which decouples estimation quality from the transmission conditions. Uncoded QPSK BER is accumulated over 2×10^6 bits per SNR point per run, across *all* 30 runs' models, with Wilson 95% confidence intervals; where zero errors are observed, the rule-of-three upper limit $3/n$ is reported instead of a point estimate of zero. SINR CDFs, sum spectral efficiency, and outage probability (threshold 0 dB) complete the link-level picture.

Electromagnetic chain. The EM experiments run on the model hierarchy of Section III-B with paired random draws across the three RIS models, at the sweep budget (4 000 training samples, 25 rounds) and the selected configuration. Every EM sweep is repeated over three data seeds, derived ratios such as the coupling-only factor are formed seed-wise before aggregation, and field-level comparisons, which share their scenes, are tested pairwise per scene (bootstrap interval and Wilcoxon signed-rank) rather than through independent standard deviations. At three seeds an interval is read as an engineering sensitivity range around large effects, and effects whose interval covers zero are reported as unresolved. The cell and the two coupling kernels are those of Sections III-B and VII-F.

TABLE 2. Full-wave models and solver settings (HFSS, $f_c = 3.5$ GHz).

Model	Geometry / sweep	Boundaries, ports	Convergence
Dipole pair	six spacings, $d/\lambda = 0.2$ to 0.7	radiation; 2 lumped	$\Delta S < 0.003$, 3–18 passes
Unit cell	two patches, $L = 14$ mm, gap 1 mm, $h = 3.048$ mm; 48 bias states $\times$ 5 angles (306 solves incl. length sweep)	periodic pairs; Floquet port, 2 modes	$\Delta S < 0.02$, 12–30 passes
Supercell	$N_c \in \{5, 7, 9\}$; six spacings	periodic + radiation; N_c lumped	$\Delta S < 0.02$, 12–30 passes
OTA testbed	32 cells, 4 BS dipoles, 2×24 probes, blocker; clean and wall variants	radiation; 84 lumped	$\Delta S < 0.02$, 6–20 passes

Full-wave simulation setup. Four full-wave models enter the paper, all solved in ANSYS HFSS with the driven modal (network) solver at $f_c = 3.5$ GHz, first-order basis functions and adaptive mesh refinement of 30% per pass (25% for the testbed), convergence being declared after two consecutive passes below the target ΔS (Table 2). Metal is modeled as a perfect conductor and the substrate as Rogers RO4003C ($\varepsilon_r = 3.55$, $\tan\delta = 0.0027$). (i) *Dipole pairs* validate the coupling kernel: two thin-wire half-wave dipoles (radius $\lambda/1000$, 0.2 mm feed gap) at six spacings, each fed by a 50 Ω lumped gap port, inside an air region whose faces lie 0.75λ from the wires and carry a radiation boundary; the solver returns Z_{11} and Z_{21}. (ii) The *unit cell* of Section VII-E is a periodic cell (period 0.5λ) with two primary/secondary pairs on its lateral faces and a two-mode Floquet port (TE_{00}, TM_{00}) 0.6λ above the patches, de-embedded to the patch plane; the varactor is a lumped series R–L–C boundary across the gap; an interpolating sweep covers $f_c \pm 4\%$ with 201 points, and incidence angles 0 to 60° in 15° steps. (iii) The *patch supercell* of Section VII-F is a finite row of N_c cells along the array axis, periodic in the transverse direction and terminated by radiation boundaries at its ends and above, with a lumped port in every varactor gap and no load, so the solver returns the open-circuit

impedance matrix of the row. (iv) The *virtual over-the-air testbed* of Section VII-J places a finite 32-element surface, a four-dipole base-station array, two tracks of probe dipoles and a conducting blocker in one radiation-bounded region, with 84 lumped ports, and returns the 84-port impedance matrix. Table 2 lists the settings that differ between the four models.

Primary contrasts and multiplicity. The paper's conclusions rest on three contrasts, which we treat as primary: the selected estimator against the deployable Kalman filter on channel NMSE over the 10×3 design, LS against the selected estimator on delivered field at the user under the scene-paired test, and FedCNN against the deployable Kalman filter on per-run BER under the paired cluster bootstrap. Every other comparison is exploratory and is read through effect sizes and intervals first, with p-values, where quoted, unadjusted. Of the primary contrasts the first two resolve with intervals clear of zero and scene-paired $p \leq 6 \times 10^{-4}$, small enough that a Holm correction across the family of three changes no conclusion; the third is reported as unresolved in Section VIII-B, which no correction alters.

TABLE 3. Simulation and training parameters (paper protocol).

Parameter	Value
BS / RIS / users / sites	$N = 16$ (ULA), $M = 32$, $K = 4$, $C = 4$
Geometry	BS $(0, 0)$, RIS $(60, 15)$ m; carrier 3.5 GHz
Numerology	$\Delta f = 15$ kHz, normal CP; $T_p = 71.35\ \mu$s; nominal slot $14T_p = 0.999$ ms
Path-loss exponents	$\alpha_{\mathrm{d}} = 3.6$, $\alpha_{\mathrm{br}} = 2.2$, $\alpha_{\mathrm{ru}} = 2.4$
Rician factors	BS–RIS 6 dB; RIS–user site-specific in $[0, 6]$ dB
Dynamics	Jakes SOS, $N_p = 16$ paths; site speeds within 3–40 km/h
Site heterogeneity	angular sectors; near 20–60 m / far 60–120 m rings; 4-point SNR sub-grids; unequal shares
Window / pilot	$T = 10$ observations; $P = M+1 = 33$ patterns each; $T_{\mathrm{obs}} = PT_p = 2.355$ ms at $M = 32$
SNR grid	-5 to 20 dB (5 dB steps); mixed-SNR training
Data	20 000 train (total) / 2 000 val / 2 000 test per SNR point
Federated schedule	$R = 50$ rounds; batch 128; weighted FedAvg; E from the playoff
Pareto PSO search	$P_s = 20$, 25 iters, $w = 0.7$, $c_1 = c_2 = 1.5$; $\lambda_C = \lambda_F = 0.02$; proxy: 5 rounds on 8k subset; 3 runs × 520 evals
Search space	$\log_{10}\eta \in [-4, -1.5]$, $H \in [16, 160]$, $E \in [1, 5]$
Warm start / hand-tuned	$(10^{-3}, 16, 3)$ / $(10^{-3}, 64, 3)$; playoff 25 rounds
Seeds	data $\{11, 22, \ldots, 110\}$ (10) × training $\{42, 123, 2024\}$ (3); EM $\{1, 2, 3\}$
Quantization	UL $\{4, 8, 16, 32\}$ bit; UL+DL $\{4, 8, 16\}$ bit; 3 seeds
Link level	pilot SNR 10 dB; QPSK; 2×10^6 bits/point/run; ZF and RZF
EM unit cell	$C_v \in [0.15, 1.5]$ pF, $R_v = 1.5\ \Omega$; phase range ≈320°; loss dip ≈3.3 dB
EM coupling / spacing	induced-EMF kernel Z_{mn}/Z_{11}; $d/\lambda \in \{0.2, \ldots, 0.7\}$; EM M-sweep $\{16, 32, 64\}$

Experiments ran in MATLAB R2026a (Deep Learning Toolbox) with ANSYS HFSS 2026 R1 for full-wave validation, on an Intel Core i5-11300H CPU (4 physical cores, 8 threads, 16 GB RAM). The studies of Sections VII-H to VII-J, the adaptation study of Section VI-D and the LMMSE comparison of Section VI-A ran on a second machine, an Intel Core i9-14900KF CPU (24 cores, 32 threads, 64 GB RAM), in MATLAB R2026a, with ANSYS Electronics Desktop 2025 R2 (HFSS) for the virtual testbed. Repeated on that machine, the column-estimation run of Section VII-D reproduces the LS and selected-estimator NMSE exactly and the FedCNN NMSE to within 0.1%.

VI. ESTIMATOR-LEVEL RESULTS

A. ESTIMATION ACCURACY

Table 4 and Supplementary Fig. S12 report the NMSE of (19) versus pilot SNR over the 10×3 cross design. All intervals quoted here are hierarchical bootstrap intervals that resample data seeds and then training seeds within them; for the learned methods the naive intervals a pooled analysis would give are 32% to 38% narrower, and are not used. Four observations follow.

First, the deployable Kalman filter is the most accurate deployable estimator in the table, at every pilot SNR. Averaged over SNR it reaches 0.294 against 0.398 for the selected estimator, and the lead holds at all six grid points. This filter estimates its transition coefficient, its signal power and its noise power from the observations it receives, so it is not an oracle; the AR(1) oracle is slightly better at 0.277 and is an upper reference rather than a competitor. We put this first because it is both the central benchmark result and the one an idealized evaluation is most likely to hide. Scoring the same estimators on observations spaced one nominal slot apart instead of PT_p, the temporal idealization of Section III-C, moves the deployable filter by 2% and the learned estimator by 45%, which is enough to reverse the ordering.

Second, least squares overtakes every learned method above 5 dB. The reversal is not specific to our estimator: at 20 dB LS (0.013) is ahead of the centralized LSTM (0.087), the selected estimator (0.088), the federated CNN (0.123) and the centralized CNN-LSTM (0.123). At high SNR the simulated observation model is exactly the linear-Gaussian one the classical estimators assume, while a network trained across the whole SNR range settles on an error floor set by its training distribution. The learned estimators lead LS only below 10 dB, and there they still trail the deployable filter.

Third, temporal structure is beneficial, but only within the learned federated class. The selected estimator attains 0.398 against 0.498 for the snapshot-based FedCNN under an identical training budget, a paired difference of −0.100 (95% CI [−0.137, −0.062], $|d| = 1.52$, win fraction 28/30), and leads its centralized LSTM counterpart by 0.043 (95% CI [−0.077, −0.016], win fraction 27/30). The first margin isolates the benefit of temporal structure within the federated class, against a snapshot federated baseline; it does not

establish superiority over classical temporal filtering. The second has the opposite sign from a privacy cost, and we therefore do not interpret it as evidence of a benefit from federation: averaging over four sites regularizes a model of this size, and regularization alone can produce an effect of this size. The defensible reading is that keeping observations at the sites had no resolvable cost here.

Fourth, the value of tuning does not resolve in this design. The untuned federated LSTM averages 0.415 against 0.398, a paired difference of -0.017 whose interval covers zero (95% CI $[-0.044, +0.012]$, $p = 0.22$, win fraction 21/30). Tuning shows in the tail rather than the mean: at 20 dB the untuned arm carries a seed spread of 0.054 against 0.026, more than double, while over the whole grid the two spreads are comparable (0.066 against 0.071). We treat that as a high-SNR effect and do not generalize it. What resolves the question is the dedicated 15-draw analysis of Section VI-B, where 93% of random configurations perform worse than the selected one; both results are kept side by side rather than only the one favoring tuning. This comparison is also the one where the runs are least correlated, with an intraclass correlation of 0.38 against 0.81 for the other two paired comparisons, giving an effective sample size of 17.1 out of 30 rather than 11.4.

Requirements of LMMSE. LMMSE is not in Table 4 because it requires the channel covariance, which no other method receives. Its limitation is quantified separately on the effective channel at 10 dB, over three data seeds and the four sites, with every sample normalized by its pilot-noise level so that one covariance applies across positions. With the population covariance of the site, LMMSE reaches 0.106, against 0.125 for LS, 0.112 for the oracle scalar Wiener filter and 0.110 for the deployable Kalman filter on the same sequences, so perfect spatial statistics are worth about as much as the temporal statistics that the Kalman filter estimates for itself. With the covariance estimated at the site from S unlabeled observations, $\hat{\mathbf{R}}_h = \hat{\mathbf{R}}_y - \sigma_e^2\mathbf{I}$ projected onto the positive semidefinite cone, LMMSE is worse than LS for $S \leq 64$ (0.289 at $S = 32$, 0.175 at $S = 64$), overtakes it at $S = 128$ (0.119) and comes within 3% of the oracle at $S = 256$ (0.109). The covariance of another site (0.109) or one pooled over all sites (0.107) serves almost as well as the site's own in this model, so the binding constraint is the number of observations, and it grows with the dimension of the estimated channel: 16 complex dimensions need 128 observations to surpass LS, whereas the cascaded channel used for configuration has $N(M+1) = 528$, with an inversion cost cubic in that dimension [7]. A spatio-temporal filter that combines the covariance with the Kalman recursion is the natural classical extension and is not evaluated here.

B. HYPERPARAMETER SELECTION: CONSISTENCY, DISCOVERY, AND RISK

The full results of the search are in Supplementary Section S8 (Supplementary Table S6 and Fig. S5). Three findings matter here. The three independent Pareto searches agree on the operating point they imply ($J = 0.4006 \pm 0.0075$), but the equal-budget random searches reach the same value (0.3990 ± 0.0077) and a front of slightly larger hypervolume (0.9285 against 0.9251), so the matched budget does not separate the two procedures. In the origin-blind playoff the hand-tuned reference wins on both J_{HF} (0.3639) and raw validation NMSE (0.3239), ahead of the three front points (0.3721 to 0.3779) and the random-search candidate (0.3790). The failure lies in the proxy rather than in the search over it: scored under the same five-round proxy that the search used, the eventual winner ranks 162nd of 228 evaluated candidates, so no search that selects by that proxy would have chosen it.

Risk reduction. Over 15 configurations drawn at random from the search space and trained under the reduced-round protocol, the untuned distribution has mean 0.324 ± 0.056 and worst case 0.455, against 0.280 for the selection; 93% of the draws are worse and 60% worse by more than 5%. Automated selection avoids the unfavorable draws rather than outperforming every favorable one, at a one-time search cost of 1263 CPU-minutes (Supplementary Section S1). The selected configuration converges to within 2% of its final validation loss by round 34 of the 50-round schedule (Supplementary Section S2).

C. ABLATIONS: TEMPORAL WINDOW, CORRELATION, AND LABELS

Two ablations at 10 dB are reported in Supplementary Section S9. Over the temporal window $T \in \{1, 2, 5, 10, 20\}$ the Kalman filters saturate by $T = 2$, while the learned estimator is best at $T = 5$ (0.167) and degrades to 0.190 at the protocol's $T = 10$, because ten observations span 23.5 ms against a 4.3 ms coherence time at 30 km/h. $T = 10$ was set under the nominal-slot interval and is kept, at a cost of about 12% to the learned estimator that changes no ordering. Over the lag-1 correlation $\rho \in \{0, \ldots, 0.99\}$, the learned estimator outperforms LS only at $\rho = 0.99$ and the deployable Kalman filter only at $\rho = 0.95$ (0.125 against 0.130), a single point rather than a regime; Section VI-D studies the same axis with retrained and conditioned networks.

D. ADAPTATION ACROSS OBSERVATION INTERVAL AND MOBILITY

With networks trained once, the cost of correcting the observation interval mixes two effects: the mismatch of fixed weights trained for another interval, and the loss of temporal information itself. To separate them, the federated LSTM was retrained and conditioned across five regimes at the main-protocol budget (20 000 training sequences, 50 rounds, three data seeds, 10 dB pilot SNR). The regimes cover the normalized Doppler axis on which the Jakes statistics depend (Table 5): the nominal slot (0.999 ms), one $M = 32$ sweep (2.355 ms, the corrected protocol), one $M = 64$ sweep (4.640 ms), and the corrected interval with every site at 40–80 km/h or at 1–5 km/h. Each regime has a specialist trained on it alone. A mixed model is trained on the union

TABLE 4. NMSE versus pilot SNR (Mean ± Std over the 10×3 seed cross design).

SNR (dB)	LS	Oracle Wiener	Kalman (oracle)	Kalman (deploy.)	CNN-LSTM (cent.) [9]	LSTM (cent.)	FedCNN [10]	FL-LSTM (unt.)	FL-LSTM (sel.)
−5	3.944	0.807	0.752	0.827	1.513 ± 0.390	1.456 ± 0.465	1.556 ± 0.289	1.226 ± 0.310	1.200 ± 0.309
0	1.238	0.559	0.499	0.522	0.599 ± 0.080	0.595 ± 0.106	0.672 ± 0.052	0.568 ± 0.089	0.571 ± 0.101
5	0.394	0.286	0.258	0.262	0.281 ± 0.019	0.266 ± 0.024	0.319 ± 0.014	0.284 ± 0.038	0.275 ± 0.030
10	0.124	0.111	0.105	0.105	0.170 ± 0.021	0.143 ± 0.013	0.183 ± 0.028	0.169 ± 0.045	0.149 ± 0.018
15	0.039	0.038	0.037	0.037	0.134 ± 0.024	0.101 ± 0.016	0.137 ± 0.032	0.129 ± 0.052	0.103 ± 0.022
20	0.013	0.012	0.012	0.012	0.123 ± 0.026	0.087 ± 0.019	0.123 ± 0.035	0.115 ± 0.054	0.088 ± 0.026
Avg.	0.959	0.302	0.277	0.294	0.470 ± 0.073	0.441 ± 0.096	0.498 ± 0.040	0.415 ± 0.066	0.398 ± 0.071

Oracle references (scalar Wiener; AR(1)-oracle Kalman) receive true statistics unavailable in deployment and are upper references, not competitors. FL-LSTM (untuned) draws one configuration at random per run. Classical columns are means over the ten data seeds; being deterministic given the data seed they carry no training-seed variation, so their intervals are t-intervals over those ten. Those intervals are uniformly narrower than the learned methods': at -5 dB the widest classical interval is ± 0.040 for LS, against ± 0.115 for the selected estimator.

of the five regimes with the same number of samples, and a conditioned mixed model also receives, as an extra input, the lag-1 correlation magnitude $|\hat{a}|$ that the deployable Kalman filter estimates from the same observations, so it can infer the regime without oracle information.

Retraining does not remove the cost. The specialist trained at the corrected interval reaches 0.155 against 0.099 for the specialist at the nominal slot, a 57% increase, while the deployable Kalman filter moves from 0.103 to 0.109 (5%) and LS is unchanged. The fixed-weight mismatch is a separate and additional cost: the slot-trained network deployed at the corrected interval yields 0.282, and 0.527 at high mobility. Conditioning removes much of that mismatch. The conditioned mixed model comes within 13 to 39% of the specialist in four regimes and is slightly better at the corrected interval (0.151 against 0.155), but no training variant improves materially on the specialist.

The study therefore narrows the claim: the learned estimator leads the deployable Kalman filter where the channel stays strongly correlated between observations, at the nominal slot (0.099 against 0.103) and at walking speed (0.061 against 0.065), and only when trained for that regime. Wherever $|\hat{a}|$ falls to 0.73 or below, every training variant trails the filter, the specialist by 42 to 53%. Correcting the protocol moves the 30 km/h operating point of an $M = 32$ surface from the first regime into the second, and that move, rather than a fragility of fixed weights, reverses the ranking.

TABLE 5. Adaptation study: NMSE at 10 dB across observation interval and mobility.

	T1	T2	T3	T4	T5
T_{obs} (ms)	0.999	2.355	4.640	2.355	2.355
Speed (km/h)	3–40	3–40	3–40	40–80	1–5
Mean $\|\hat{a}\|$	0.94	0.73	0.53	0.39	0.99
LS	0.125	0.124	0.126	0.124	0.124
Kalman (dep.)	0.103	0.109	0.108	0.107	0.065
LSTM, specialist	**0.099**	0.155	0.160	0.164	**0.061**
LSTM, trained on T1	**0.099**	0.282	0.456	0.527	0.083
LSTM, mixed	0.134	0.162	0.194	0.238	0.114
LSTM, mixed + $\|\hat{a}\|$	0.124	0.151	0.181	0.196	0.085

Means over three data seeds. T2 is the corrected protocol of Table 3; T1 the nominal slot. Bold: learned estimator ahead of the deployable Kalman filter.

E. LABEL SOURCE: WHAT REALISTIC ACQUISITION COSTS

Repeating the main comparison with teacher labels from an occasional long-pilot acquisition ($R_t = 4$ repetitions of the sweep at +10 dB, followed by the same cascaded LS) instead of genie labels, with the same seeds and test sets, changes the SNR-averaged NMSE of FedCNN by −0.2% and of the selected estimator by +0.2%, neither distinguishable from zero (paired bootstrap $p = 0.64$ and 0.88); the classical estimators are label-free and do not change. The table is in Supplementary Section S3.

F. ROBUSTNESS TO SYSTEM SCALE

Two scale sweeps are reported in Supplementary Section S4. Across $M \in \{16, 32, 64, 128\}$, which also stretches T_{obs} from 1.21 to 9.20 ms, LS degrades gently (0.117 to 0.141 at 10 dB) while the selected estimator degrades steeply (0.109 to 0.201) and leads LS only at $M = 16$; it stays 2 to 35% below FedCNN at every size. The M-sweep therefore quantifies the joint cost of surface size and pilot length, not an aperture effect. Across federations of $C = 4$ to 32 sites with fixed per-site data, NMSE improves up to $C = 8$ (0.160 to 0.139) and rises again beyond it, the margin over FedCNN stays between 5 and 16%, and half participation is worse than quarter participation at three of the four sizes, a reproducible feature of this schedule that remains unexplained.

VII. ELECTROMAGNETIC AND PHYSICAL VALIDATION

A. ELECTROMAGNETIC CONSISTENCY: UNIT CELL, COUPLING KERNEL, AND VALIDATION

Supplementary Fig. S9 characterizes the electromagnetic building blocks. Panel (a) shows the unit-cell $|\Gamma|$ versus frequency for five bias states: shifting the varactor bias moves the resonance through f_c, and the response is visibly dispersive. Panel (b) is the resulting S-curve at f_c: the achievable phase range is 319.6° with a minimum amplitude of 0.68 (−3.3 dB) at resonance, which is the amplitude–phase coupling and loss expected of a lossy resonant varactor-loaded cell and absent by construction from the ideal phase-shifter abstraction. Panel (c) shows the TE/TM reflection phase versus incidence angle up to 60°: the two polarizations separate further with obliquity, another effect the idealized model omits. Panel (d) plots the coupling kernel $|K| = |Z_{21}/Z_{11}|$ against element spacing, growing from 0.29 at $d = 0.7\lambda$ to 0.65 at $d = 0.2\lambda$.

The kernel itself is validated three ways. First, the closed forms reproduce the classical values $Z_{11} = 73.1 + j42.5\,\Omega$ and $Z_{21}(0.5\lambda) = -12.5 - j29.9\,\Omega$. Second, independent

numerical quadrature of the sinusoidal-current near-field EMF integral agrees with the closed forms to machine precision ($\sim 10^{-16}$ relative), as does the algebraic identity between the Foldy–Lax form (7) and the multiport solve (8). Third, full-wave HFSS two-port simulations of thin-wire dipole pairs across $d/\lambda \in [0.2, 0.7]$ (adaptive solves at f_c, $\Delta S < 0.003$) agree with the analytic kernel in the quantity the generator uses, $K = Z_{21}/Z_{11}$: the magnitude matches within 0.7–4.6% (mean 2.7%) over the full spacing range, with a systematic $\approx 10^\circ$ phase offset attributable to the feed-gap reactance entering the full-wave Z_{11} reference. Raw Z_{21} agrees up to a single spacing-independent complex factor ($\approx 22\%$ when compared directly, uniform across d), the classical signature of terminal-current versus sinusoidal-basis referencing (Supplementary Fig. S7). The coupling model's spacing profile is therefore checked against Maxwell-solver ground truth rather than assumed.

B. IMPACT OF MUTUAL COUPLING ON CHANNEL ESTIMATION

Fig. 2 sweeps the element spacing d/λ from 0.7 down to 0.2 under the three-model hierarchy with paired random draws. Three observations follow. First, the unit-cell response alone (EM-nc: amplitude–phase coupling, loss, 320° phase clipping) costs a modest, spacing-independent penalty: 8.9 to 10.7% in NMSE for LS at every spacing, 9.4 to 12.7% for the deployable Kalman filter, and 5.3 to 9.0% for the two learned estimators. Second, mutual coupling degrades every estimator substantially, and monotonically once d falls below $\lambda/2$: the LS pilot estimator rises from 0.172 at $d = 0.5\lambda$ to 0.229 at 0.4λ, 0.325 at 0.25λ and 0.359 at 0.2λ. Isolated from the cell response (EM-mc relative to EM-nc, seed-paired), the coupling factor grows from 1.16 to 1.29× at 0.7λ to 1.96 to 2.73× at 0.2λ depending on the estimator (1.27× and 2.66× for LS). The two learned methods carry the *smallest* coupling factors at every spacing (1.96 and 2.10× against 2.66 and 2.73× at 0.2λ), so the robustness that learning provides against coupling is real but starts from a higher absolute error and does not close the gap. Third, the degradation does *not* reorder the estimators. In the seed means the ordering is the same at all six spacings: the deployable Kalman filter first (0.149 at 0.5λ rising to 0.324 at 0.2λ), then LS, then the selected estimator, then FedCNN. With three seeds the marginal intervals of the two learned methods overlap, so the means alone are inconclusive; because the seeds are paired across spacings, the claim is instead checked and stated cell by cell. Over the 18 seed-by-spacing cells the deployable filter leads LS, the selected estimator and FedCNN in all 18, and LS leads FedCNN in all 18. LS leads the selected estimator in 16 and the selected estimator leads FedCNN in 16, so the complete four-way ordering holds in 14 of the 18 cells; each of the four exceptions is a margin below 0.007 in NMSE, and the seed-wise mean over the six spacings has the stated sign for all six pairwise relations in every seed. The classical lead is unqualified: no learned estimator leads the deployable filter at any seed or spacing. Coupling therefore acts as a common degradation that every estimator inherits in similar proportion, not as a systematic bias that only label-driven learning can correct.

Supplying the multiple-scattering inverse with the surface reflection coefficient of (5) instead of the Z_{11}-referenced quantity (6) implies negative element resistance over 91–95% of the bias range. Re-running the spacing sweep under that convention with only the reference changed (Table 6) inflates the LS NMSE 2.3× at 0.7λ and $96.4 \pm 19.1\times$ at 0.2λ, against $2.9 \pm 0.44\times$ for the learned estimator, which is retrained on the defective generator and adapts to it, whereas training-free LS inherits it unaltered; $\|\mathbf{\Phi}_{\text{eff}}\|_2$ rises from at most 1.53 to 893. The artifact is numerically silent. Any study using (7) can rule it out by checking $\Re Z_{\text{L},m} \geq 0$ for $Z_{\text{L},m} = Z_{11}(1/\Phi_m - 1)$, which our generator does for every dataset (full discussion in Supplementary Section S18).

Under consistent scaling at $d = 0.5\lambda$ the NMSE grows with M for both LS (0.157 to 0.214 from $M = 16$ to 64) and the selected estimator (0.190 to 0.274), LS leads at every size by 21 to 28%, and at $M = 64$ the sweep no longer fits the coherence interval at the reference speed; replacing spherical by far-field propagation moves the NMSE by less than 0.2% at $M = 32$ (Supplementary Section S19).

TABLE 6. Impedance-reference ablation: NMSE at 10 dB under EM-mc with the element response referred to Z_{11} (Correct) versus to $\eta_0/\cos\theta_i$ (Mis-referenced), over three independent data seeds.

	LS		Selected		Inflation	
d/λ	corr.	mis-ref.	corr.	mis-ref.	LS (×)	sel. (×)
0.20	0.367	35.341	0.383	1.103	96.4 ± 19.1	2.88 ± 0.44
0.25	0.315	3.452	0.349	0.988	10.9 ± 0.5	2.83 ± 0.38
0.30	0.284	1.286	0.324	0.883	4.5 ± 0.4	2.73 ± 0.36
0.40	0.227	0.993	0.275	0.758	4.4 ± 0.4	2.76 ± 0.61
0.50	0.179	1.154	0.233	0.715	6.4 ± 1.6	3.11 ± 1.14
0.70	0.177	0.409	0.231	0.425	2.3 ± 0.3	1.85 ± 0.30

Identical seeds, estimators and protocol within the pair; only the impedance reference of the element response inside (7) differs. The mis-referenced column is not a valid physical model, since it implies negative element resistance over most of the bias range, and is reported solely to quantify the artifact. The defective convention is retained in the implementation as a documented switch defaulting to the correct reference, so the artifact can be re-created at will rather than existing only in this table. Inflation ratios are formed seed-wise and averaged afterwards, so the pairing survives aggregation; ± figures are 95% intervals over the three seeds. The absolute NMSE columns run under the reduced sweep budget and should not be read cell-by-cell against Section VII-B; what transfers is the ratio.

C. ELECTROMAGNETIC MODEL MISMATCH

Training on the ideal model and deploying under the full EM model, with paired draws, costs the selected estimator +2.9% at $d = 0.5\lambda$ and +7.6% at 0.25λ (FedCNN +3.9 and +6.8%), and the reverse direction costs at most a few percent; with three seeds, seven of the eight relative costs have intervals that cover zero. Whatever a network learns about coupling is therefore worth at most a few percent in either direction at these spacings, too little to support the claim that a learned estimator implicitly corrects the physics (Supplementary Section S5).

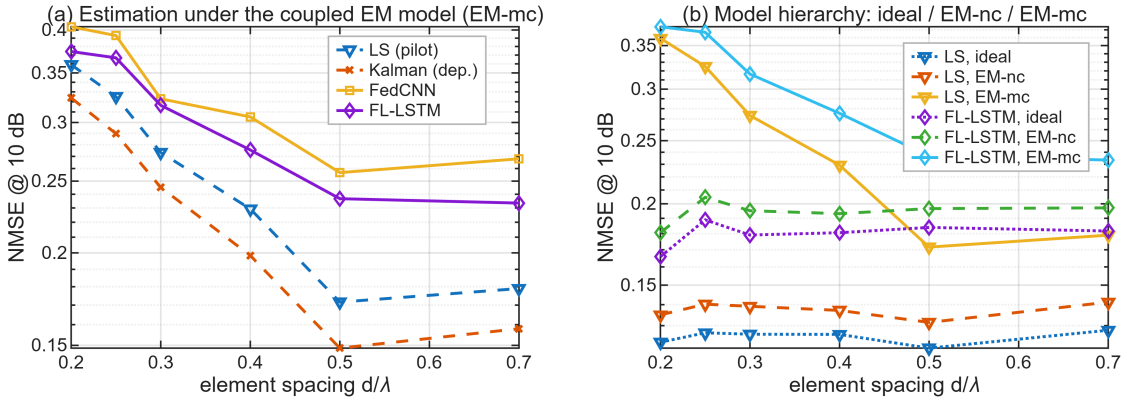


FIGURE 2. **(a) NMSE at 10 dB versus element spacing d/λ for all estimators under the full EM model. (b) The ideal/EM-nc/EM-mc model hierarchy over the same spacings, for LS and the selected estimator, on paired random draws.**

D. FIELD-LEVEL VALIDATION: CURRENTS, PATTERNS, AND FOCUSING

NMSE is a channel-domain metric; Supplementary Fig. S8 tests whether estimation gains carry over to actual wave control. For each of 50 random user scenes under the full EM model (at both the default $d = 0.5\lambda$ and the dense $d = 0.25\lambda$), the per-element cascade columns are estimated (LS directly; FedCNN and the selected FL-LSTM as federated column denoisers with identical hyperparameters), the RIS bias states are configured from each estimate through the calibrated lookup table, the realized coupled surface response (7) determines the element currents, and the scattered field $|E(x, y)|$ is computed over the deployment area. Perfect-CSI configuration provides the reference. The chain is thus exactly the deployed one: estimate → bias commands → coupled currents → field.

The scene-level intervals and p-values below condition on the trained models: the 50 scenes are paired draws of the user geometry, not replicates of training, so method-level inference rests on the sign of the paired difference across three training seeds. LS delivers more field at the user than the selected estimator under all three seeds, by 0.36 to 0.42 dB at $d = 0.5\lambda$ and 0.47 to 0.72 dB at 0.25λ, while the per-column NMSE lead spans 8.9 to 10.4× and 6.8 to 7.3×. The full statement of this scope is in Supplementary Section S15.

The column-estimation separation is wide: at $d = 0.5\lambda$ the column NMSE is 1.088 for the selected estimator against 1.003 for FedCNN and 10.560 for LS, a 9.7× lead over LS (7.3× at 0.25λ). Neither the surface currents nor the radiated field inherits that ordering. The surface-current error is unresolved at 0.5λ (0.349 for the selected estimator against 0.330 for LS, Wilcoxon $p = 0.40$) and resolved against the learned estimator at 0.25λ (0.852 against 0.677, $p = 2 \times 10^{-5}$). LS delivers more field at the user at both spacings, 0.36 dB (interval $[0.11, 0.65]$, $p = 6 \times 10^{-4}$) and 0.47 dB ($[0.16, 0.75]$, $p = 1 \times 10^{-4}$); peak realized gain and sidelobe level favor LS as well, four of five field metrics at 0.25λ and three of five at 0.5λ are resolved deficits for the learned estimator, none is a resolved gain, and the pointing error is unresolved (Supplementary Section S14). Configuration tolerates column error well: even configurations computed from the noise-dominated LS columns focus within 0.89 dB of perfect CSI at 0.5λ, against 1.25 dB for the selected estimator. Section VII-I traces the reversal to the error components that phase-only configuration ignores and NMSE counts in full. The absolute differences are modest, 0.36 to 0.68 dB between the best and the worst configured surface against 0.89 to 2.20 dB to perfect CSI, but the result cautions against reading NMSE ratios as beamforming ratios.

E. FULL-WAVE CHARACTERIZATION OF THE UNIT CELL

Every result above rests on the equivalent-circuit unit cell of Section III-B, whose reflection coefficient follows in closed form. That model is standard but approximate. Here it is replaced by the computed response of a geometry-resolved varactor-loaded unit-cell model, a full-wave HFSS characterization under Floquet periodicity, to see what the idealization costs at channel level.

Cell and solver setup. Two rectangular patches sit on a grounded 3.048 mm RO4003C substrate ($\varepsilon_r = 3.55$, $\tan\delta = 0.0027$); the central gap is bridged by a series R–L–C varactor ($R_s = 1.5\,\Omega$, package $L_s = 0.45$ nH, $C_v = 0.15$–1.5 pF). Scan-angle primary/secondary pairs close the four lateral faces and a two-mode Floquet port (TE_{00}, TM_{00}) terminates the cell from above, so both polarizations come from a single solve; the reference plane is de-embedded to the patch plane. The period equals the array element spacing (0.5λ), which keeps the cell consistent with the coupling kernel of (7).

The patch length ($L = 14$ mm) is selected by a full-wave parametric sweep over $L = 5$ to 20 mm, the bias axis is refined until the worst phase step falls to 12.4° at normal incidence (48 bias states × 5 angles), every solve is required to run at least 12 adaptive passes to avoid a spurious early convergence on the coarse mesh, and coupling is counted once per arm, the embedded Floquet response being used with the kernel disabled because de-embedding it with the dipole row sum returns non-passive values. These details are given in Supplementary Section S11.

The full-wave cell model exhibits a narrower phase range and a shallower resonant dip. At f_c under normal incidence (Fig. 3), the full-wave cell spans 251.6° of phase against 319.6° for the equivalent-circuit model, while its worst-case insertion loss is *lower* (2.07 against 3.34 dB). Its *mean* loss is the higher of the two, however (1.23 against 1.07 dB), so the full-wave cell is not uniformly less lossy; its response is flatter, trading a shallower resonance for more loss away from it. The circuit model is thus optimistic about attainable phase coverage, pessimistic about resonant loss and optimistic about loss over the rest of the bias range; these errors act on different parts of that range and do not cancel.

Why the two curves disagree, and a calibrated circuit. The circuit of (4) is a two-branch sheet model, and the geometry solved in HFSS has three properties that it leaves out. The first is the fringing capacitance across the varactor gap and between the patch edges, which sits in parallel with the varactor branch, lowers the slope of the phase response and narrows its swing. The second is the transformation from a lumped load to a sheet impedance: the voltage across a 1 mm gap is not the mean tangential field over a 42.8 mm cell, so the varactor branch enters Z_s through a scale factor rather than one to one. The third is anisotropy. The varactor bridges the gap along one axis only, so the cell is tuned for the field across the gap and is almost inert for the orthogonal field, whose phase moves by 0.3° over the whole bias range in HFSS, whereas (5) applies one Z_s to both polarizations. The way a lumped device is attached also changes the response. For a PIN-diode cell designed at 29 GHz, modeling the diode as an edge or a face lumped element instead of a three-dimensional diode model moved the reflection band of one bias state from 29 GHz to about 19 or 23 GHz, and the shift differed between bias states [41]. To test whether these properties account for the gap, the circuit was extended by a shunt capacitance C_p, a shunt loss R_p, a scale factor κ on the varactor branch and a package inductance L_s,

$$Z_s^{-1} = \frac{1}{j\omega L_p} + j\omega C_p + \frac{1}{R_p} + \frac{1}{\kappa\left(R_v + j\omega L_s + 1/(j\omega C_v(u))\right)}, \tag{23}$$

which reduces to (4) for $C_p = 0$, $R_p \to \infty$, $\kappa = 1$, $L_s = 0$, and fitted together with L_p, R_v and a reference-plane phase to the 48 HFSS bias states at f_c under normal incidence (multi-start Nelder–Mead on log-parameters, complex reflection error). The fit lowers the RMS phase error from 82.7° to 1.0° (worst state 2.2°) and the RMS error in $|\Gamma|$ from 0.98 to 0.02 dB, and it reproduces the phase range (251.3° against 251.6°), the depth of the loss dip (−2.06 against −2.07 dB) and its bias position ($u = 0.68$ in both, against 0.51 for (4)). The fitted values are physically plausible: $R_v = 1.46\,\Omega$ against the 1.5 Ω assigned to the varactor in HFSS, $L_p = 2.90$ nH against 2.51 nH, $C_p = 0.41$ pF, $\kappa = 3.57$, and $R_p = 21\,\text{k}\Omega$, i.e. negligible shunt loss. The disagreement in Fig. 3(a) and (b) is therefore a missing-element error of the circuit of (4) rather than a solver artifact, and the dotted curves show the calibrated circuit closing it. Panel (d) shows the anisotropy directly. For the tuned polarization the calibrated circuit stays within 2° of HFSS up to 30° incidence and within 7.3° at 45°, and departs at 60° (9.6° against −23.8°), where the patch cell is spatially dispersive and a local sheet model no longer applies. The orthogonal polarization stays between −173° and −197° with $|\Gamma|$ within 0.03 dB of unity, the response of an unloaded patch over ground. The calibrated cell is carried through the full estimation chain as a fifth cell family in Section VII-H.

The component-level gap has no resolvable channel-level cost. Retrained on each cell and cross-evaluated over nine operating conditions (three incidence angles × three frequency offsets), a model trained on the circuit cell and deployed on the full-wave cell stays within the seed spread of one trained on the full-wave cell itself (0.338 ± 0.016 against 0.353 ± 0.015 at f_c; the circuit-trained model leads in eight of nine conditions by a mean of 0.007). The circuit is therefore an adequate training surrogate even where it is not an accurate description (Supplementary Section S20).

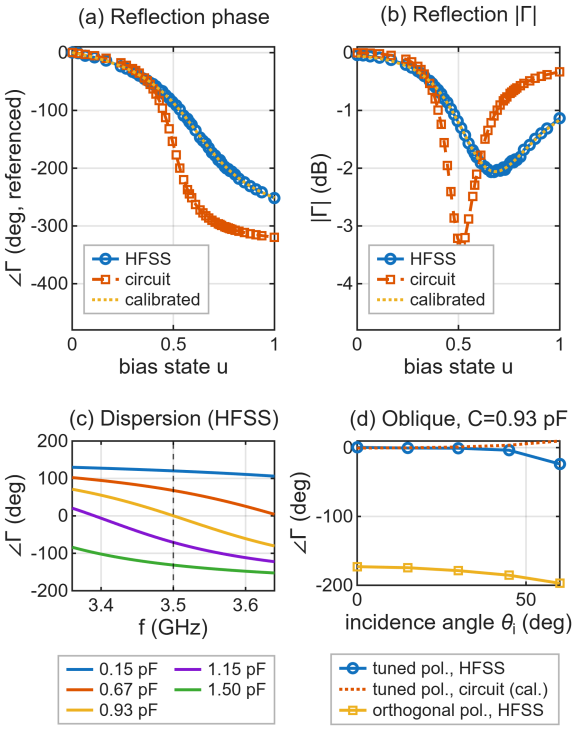


FIGURE 3. Full-wave unit cell at f_c. (a) Referenced phase and (b) $|\Gamma|$ versus bias: HFSS, circuit (4) and calibrated circuit. (c) HFSS dispersion. (d) Phase versus incidence angle at 0.93 pF.

F. A COUPLING KERNEL FROM THE PATCH GEOMETRY

The chain has so far combined two inconsistent descriptions. The cell response comes from a geometry-resolved varactor-loaded patch model characterized in HFSS, while the element interaction comes from induced-EMF formulas for a canonical thin half-wave dipole. Both are defensible in isolation but do not describe the same object, which is why the full-wave

arm of Section VII-E ran with the analytic kernel disabled: an embedded Floquet response already accounts for the neighboring elements. That comparison therefore validates the cell alone, not the cell together with its coupling.

The inconsistency is resolved by extracting the kernel from the same geometry. In the multiport form $\mathbf{Z}_{\mathrm{RIS}}$ is the open-circuit impedance matrix of the array with every load removed, which has a direct computational counterpart: a finite supercell of patch cells, periodic transverse to the array, radiating at the ends, with a lumped port at each varactor gap and no RLC load. The solver returns $\mathbf{Z}_{\mathrm{RIS}}$ itself, so

$$K_{mn}^{\mathrm{patch}} = Z_{mn}^{\mathrm{patch}} / Z_{11}^{\mathrm{patch}} \tag{24}$$

follows without any further assumption, and the varactor stays where it belongs, in diag(Z_s). One solve per spacing yields every neighbor order at once, and averaging same-order pairs over the interior cells suppresses the edge effect of a finite supercell. Repeating the extraction at $N_c = 5$, 7 and 9 with the geometry otherwise fixed bounds the residual truncation error: going from the $N_c = 7$ used here to $N_c = 9$ moves K_1 by 0.05% at $d = 0.25\lambda$ and 0.65% at 0.5λ, and moves K_1, K_2 and Z_{11} by at most 1.01% at either spacing, whereas the step from $N_c = 5$ to 7 reaches 8.8%. The test is therefore sensitive enough to detect an unconverged supercell, and $N_c = 7$ is converged.

TABLE 7. Patch-derived coupling kernel against the canonical dipole kernel.

d/λ	$\lvert K_1^{\mathrm{p}}\rvert$	$\lvert K_1^{\mathrm{d}}\rvert$	ratio	$\Delta\angle K_1$	Z_{11}^{patch} (Ω)
0.20	0.255	0.648	0.393	24.8°	$9.4 + j48.8$
0.25	0.195	0.587	0.333	28.8°	$9.6 + j46.2$
0.30	0.154	0.534	0.288	31.6°	$9.8 + j43.7$
0.40	0.102	0.448	0.228	33.2°	$10.1 + j39.1$
0.50	0.077	0.384	0.200	28.6°	$10.8 + j35.7$
0.70	0.055	0.294	0.188	9.0°	$12.2 + j29.4$

First-neighbor entries; the multiport solve yields the higher orders as well. Superscripts p and d denote the patch-derived and dipole kernels. Z_{11}^{patch} is the interior-cell self-impedance returned by the multiport solve with all loads removed, averaged over interior cells of the $N_c = 7$ supercell. The patch kernel is weaker than the canonical dipole kernel at every spacing, by a factor between 2.5 and 5.3, and the two also disagree in phase.

Table 7 compares (24) against the dipole kernel in magnitude and phase across the swept spacings. Repeating the spacing sweep of Section VII-B with the patch kernel then answers the question that inconsistency raised: whether the coupling conclusions depend on which kernel was used or only on the presence of coupling.

The conclusions hold, but the magnitude changes. With the same seeds, estimators and protocol, and only the kernel changed, the coupling-only factor at $d = 0.2\lambda$ falls from 1.96–2.73× across the four estimators to 1.11–1.19×, so the excess penalty that coupling imposes shrinks by roughly an order of magnitude. The spacing dependence largely disappears with it: under the patch kernel LS moves only from 0.135 at 0.5λ to 0.159 at 0.2λ, against 0.172 to 0.359 under the dipole kernel. What does not change is the ranking. Under both kernels the seed-mean ordering at all six spacings is the deployable Kalman filter, then LS, then the selected estimator, then FedCNN. Checked cell by cell over the three seeds as in Section VII-B, the patch kernel gives fewer exceptions: five of the six pairwise relations hold in all 18 seed-by-spacing cells and the complete four-way ordering in 16, against 14 under the dipole kernel, with the exceptions again margins below 0.01 in NMSE. Since the patch kernel and the cell describe one geometry, we read its numbers as the physically consistent account of what coupling costs and the dipole numbers as a conservative strong-coupling stress case for this cell. The practical reading is that the canonical dipole kernel overstates how much dense packing costs a patch array of this geometry, while the estimator ordering is a property of the estimators and not of either kernel.

G. A TWO-DIMENSIONAL APERTURE

On an 8 × 8 aperture whose kernel comes from the generalized induced-EMF integral, with transverse, axial and staggered neighbors from one expression (Supplementary Section S10), the conclusions reached on the line hold. The deployable Kalman filter is first at every model level and the selected estimator trails LS in channel NMSE under coupling (0.326 against 0.234). It leads LS by 4.0× in per-column NMSE (1.82 against 7.25), yet delivers 0.40 dB less field at the user (interval $[-0.80, -0.15]$, $p_W = 2 \times 10^{-4}$) and a larger surface-current error, with the sign of the field difference stable over three training seeds. The planar kernel keeps $\Re\,\mathbf{Z}_{\mathrm{RIS}}$ positive semidefinite, but its smallest eigenvalue falls from 53.6 on the line to 0.123, so denser or larger planar apertures should be checked for conditioning before the multiport inverse is trusted.

H. GENERALITY ACROSS UNIT-CELL FAMILIES

The impedance-reference audit and the estimator ordering were established above on one cell family, and the size of the audit differs between the two kernels. To separate what is general from what is specific to that cell, the spacing sweep was repeated on five cell families, under both kernels and both references, with paired seeds (three data seeds, $d \in \{0.5, 0.25, 0.2\}\lambda$, $M = 32$, 10 dB pilot SNR, the sweep budget of Table 6). Besides the baseline cell, the families change one property at a time: a low-loss cell ($R_v = 0.5\,\Omega$), a high-loss cell ($R_v = 4\,\Omega$) and a narrow-tuning cell (C_v = 0.30–1.00 pF), each with L_p re-tuned for mid-bias resonance at f_c, and the HFSS-calibrated circuit of (23). Together they span phase ranges from 251° to 320° and resonant losses from 1.1 to 9.7 dB (Table 8).

Three conclusions hold in every family. First, under the self-impedance reference the deployable Kalman filter is the most accurate estimator in all 135 seed-level cases (five families, three model levels, three spacings, three seeds), the seed-mean ordering (Kalman, then LS, then the learned estimators) holds in all 45 family-level cases, and LS outperforms both learned estimators in 131 of the 135 seed-level cases, the four exceptions being small margins in the low-loss

and calibrated families. Second, the mis-referenced model inflates LS more than the learned estimator in every family and under both kernels, so the direction in which the artifact favors learning does not depend on the cell. Third, under the dipole kernel that bias reverses the ranking in every family at $d \leq 0.25\lambda$, where a learned estimator becomes the most accurate one.

The magnitude does not generalize. At 0.2λ the LS inflation under the dipole kernel ranges from $9.3\times$ (high loss) to $87\times$ (baseline), against 2.3 to $3.1\times$ for the selected estimator; it tracks how close ($\mathbf{I} + \mathbf{\Phi K}$) comes to singularity, which depends on the loss and the tuning range of the cell. The baseline value under these seeds, $87\times$, lies inside the $96 \pm 19\times$ of Table 6. Under the patch kernel the same audit gives 1.4 to $2.1\times$ for LS and 1.2 to $1.7\times$ for the selected estimator: the direction is the same, but the weaker kernel never brings the inverse close enough to singularity to reverse the ranking. The 96-fold figure is therefore a property of the dipole kernel and the baseline cell, and what transfers is the direction of the artifact together with its dependence on kernel strength. The HFSS-calibrated family behaves like the others; its higher absolute NMSE follows from its narrower phase range, which clips more of the commanded phases (Supplementary Section S27).

TABLE 8. Generality across cell families: impedance-reference inflation at $d = 0.2\lambda$.

Family	Change	Range	Dip	Inflation, LS / sel. (×) Dipole	Patch
Baseline	—	319.6°	3.34 dB	87.1 / 2.94	1.82 / 1.49
Low loss	$R_v = 0.5\,\Omega$	319.5°	1.10 dB	14.8 / 2.29	2.12 / 1.65
High loss	$R_v = 4\,\Omega$	320.3°	9.70 dB	9.3 / 3.07	1.36 / 1.22
Narrow	C_v 0.3–1.0 pF	284.1°	2.08 dB	28.5 / 2.64	1.79 / 1.47
Calibrated	(23)	251.3°	2.06 dB	20.8 / 2.80	1.50 / 1.31

Ratio of mis-referenced to correctly referenced NMSE, formed seed-wise and averaged over three seeds. Range and dip: phase range and resonant loss at f_c, normal incidence.

I. WHY CHANNEL NMSE MISRANKS THE DELIVERED FIELD

Section VII-D found that a near-tenfold lead in column NMSE delivers less field. This subsection quantifies the mechanism on the same chain, over 200 scenes (the first 50 are those of Section VII-D) and three training seeds, in four RIS configurations: the linear array at 0.5λ and 0.25λ, the 8×8 aperture of Section VII-G, and the linear array with the HFSS-calibrated cell of (23). Five beamforming and RIS-control strategies are applied in each: the SVD beam of Section VII-D, uniform BS weights, joint alternating optimization of BS maximum-ratio transmission and RIS phases driven by the same estimate, and the SVD beam with 2-bit and with 1-bit RIS phase quantization. Two field measures are reported, the free-space field at the user as in Section VII-D and the amplitude received through the modeled cascade $\mathbf{g}_{\mathrm{ru}}$, which also contains its NLoS part.

The mismatch persists everywhere. LS delivers more field at the user than the selected estimator in all 60 combinations of configuration, strategy and training seed, by 0.14 to 0.95 dB (0.19 to 0.76 dB through the cascade). The alternating optimization converges to the SVD beam, because the near-field BS–RIS channel is close to rank one, so its numbers coincide with those of the SVD strategy. Three computed diagnostics explain the gap, and they act in sequence.

Error covariance across the BS antennas. Configuration does not use an estimated column directly. It projects the column on the calibrated $\mathbf{H}_{\mathrm{br}}$ column, an N-antenna matched filter, and only the error component along that column survives. The share of the column-error energy along it is 0.25 for LS against 0.64 for the selected estimator at 0.5λ, and 0.20–0.37 against 0.63–0.77 over the four configurations; white error would give $1/N = 0.06$. The LS error is largely pilot noise spread over the antennas, whereas the learned residual lies along the channel it estimates. The projection therefore removes 89% of the LS error energy and 30% of the learned one, and a $6.6\times$ lead in scene-mean column NMSE ($9.7\times$ on the pooled test set) becomes a $1.06\times$ lead in the projected element coefficients (0.314 against 0.297). A controlled test isolates this step (Fig. 4(d)): at the selected estimator's own column NMSE, error white over the antennas costs 0.11 dB at the user, and the same error energy aligned with $\mathbf{H}_{\mathrm{br}}$ costs 1.69 dB.

Amplitude and phase. A phase-only surface uses only the phase of each element coefficient. Of the selected estimator's element NMSE, 54% is amplitude error (44% for LS), and in the controlled test an amplitude-only error at that element NMSE costs 0.00 dB against 1.47 dB for a phase-only error of the same NMSE.

Error-phase coherence. A phase error common to all elements rotates the whole aperture and costs almost nothing (0.03 dB in the test); the spread around it reduces the field. The relevant statistic is the error-phase coherence $\rho = |\sum_m w_m e^{j\delta_m}| / \sum_m w_m$, with $\delta_m = \angle(\hat{g}_m / g_m)$ the phase error of the estimated element coefficient $\hat{g}_m$ against its perfect-CSI value g_m and w_m the contribution of element m to the perfect-CSI field; ρ is the field ratio of an ideal uncoupled phase-only surface. Like NMSE, ρ needs the true channel: it is an offline evaluation diagnostic, not a quantity an estimator can compute in operation. LS carries the larger common-mode error (14.0° against 5.8°) and the smaller spread (23.5° against 27.7°), hence the higher coherence (0.905 against 0.878). Evaluated before the bias table, loss and coupling act, $20\log_{10}\rho$ correlates with the realized loss through the cascade at $r = 0.95$ to 0.97 per scene across the four configurations, against 0.81 to 0.90 for element NMSE and 0.39 to 0.64 for column NMSE (Fig. 4(a), (b)). Scene by scene, the sign of the LS-minus-learned field difference agrees with the sign of the spread difference in 72 to 84% of scenes and with that of the column-NMSE difference in only 24 to 33%, so column NMSE misranks the two estimators in most scenes.

Beamforming sensitivity. A phase-only perturbation sweep through the full chain (Fig. 4(c)) follows $10\log_{10} e^{-\sigma^2}$, with σ the phase spread in radians, the large-aperture limit for Gaussian phase error, to within 0.1 dB from 5° to 60° of

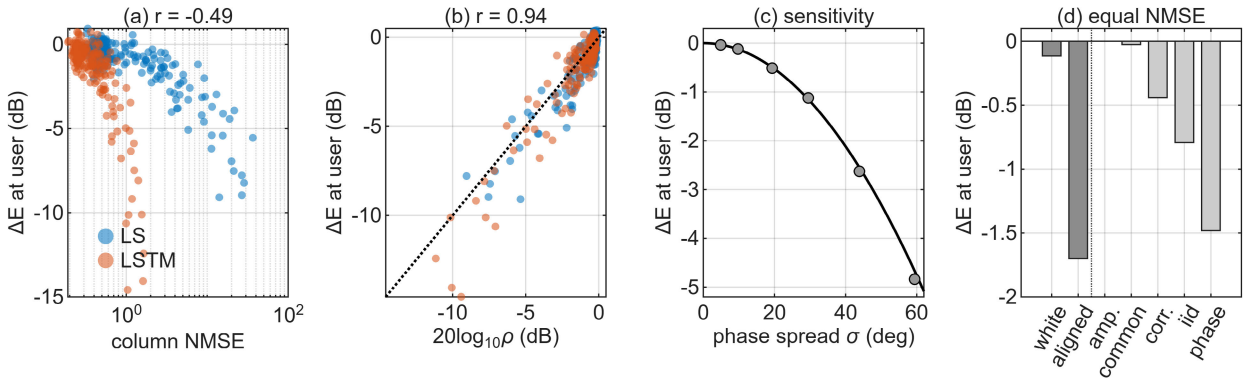


FIGURE 4. Mechanism of the NMSE-to-field mismatch ($d = 0.5\lambda$, SVD beam). (a), (b) Per-scene field loss at the user against column NMSE and against $20\log_{10}\rho$. (c) Phase-only sweep (markers) and $10\log_{10} e^{-\sigma^2}$ (line). (d) Controlled errors at equal NMSE.

spread; near the operating point the 4° difference in spread is worth 0.3 dB, the size of the simulated field gap.

What the controlled tests rule out. A residual correlated across elements is not the mechanism. The learned residual is more correlated than the LS residual only in amplitude; the cross-element correlation of the phase errors is the same for both (0.26, null 0.12 to 0.14), and at equal element NMSE correlated error costs *less* than i.i.d. error (0.44 against 0.78 dB) because it is partly common-mode (Supplementary Section S26). The resulting mechanistic chain is therefore the following: the learned residual lies along the channel and survives the matched-filter projection; of what survives, the amplitude part costs nothing and the phase part costs field through its within-aperture spread; channel NMSE gives full weight to the components that configuration ignores.

J. A VIRTUAL OVER-THE-AIR TESTBED

A fabricated surface and an over-the-air campaign were not available for this work. An end-to-end estimation experiment needs a varactor-loaded board with an independent bias line for each of 32 diodes, a controller synchronized with the pilot sweep at the OFDM symbol rate, a multi-antenna base station and an anechoic or controlled-multipath chamber. Instead, a controlled full-wave virtual testbed tests whether the conclusions persist beyond the analytic propagation model: a laboratory-scale virtual scene is solved in HFSS as one network (Table 2, Fig. 5(a)), comprising a finite board of 32 cells of the type of Section VII-E with a lumped port in every varactor gap, a four-dipole base-station array, two tracks of 24 probes spaced by the distance a user covers in one observation interval (20 mm per 2.355 ms, i.e., 30.6 km/h), and a metal blocker that removes the line-of-sight path. One solve (14 adaptive passes, $\Delta S = 0.011$, 658 604 tetrahedra) returns the 84-port impedance matrix. From it the transmission between base-station and probe ports, the quantity an over-the-air measurement records, follows for any bias state by network theory without re-solving (Supplementary Section S33). A second solve replaces one lateral radiation boundary by a metal wall, adding strong specular multipath (14 passes, $\Delta S = 0.011$, 617 837 tetrahedra); the wall halves the configuration-dependent share of the received power, from 0.31 to 0.14.

Model checks. Both solves are reciprocal to 10^{-16} and passive. The board's gap-port self-impedance, $13.3 + j36.3\ \Omega$ for an interior cell, agrees with the supercell value of Table 7 ($10.8 + j35.7\ \Omega$), and its first-neighbor coupling, $|K_1| = 0.096$, agrees with the patch kernel (0.077) rather than with the dipole kernel (0.384). With the full-wave propagation held fixed and only the RIS model swapped (Fig. 5(b)), the self-referenced form with the full-wave kernel is the exact network-theoretic reduction of the exported full-wave multiport model; it matches the direct network solution to 5×10^{-16}. With the patch kernel of Section VII-F it reproduces the configuration-dependent part of the channel within 5.0%, against 23.8% without coupling, 22.3% when the surface-referenced reflection enters the inverse, and 66.9% with the dipole kernel, which overstates the coupling of this cell and is less accurate than omitting it; with the wall the four errors are 4.8, 32.2, 28.9 and 68.1%. The impedance reference therefore matters in the full-wave scene even at half-wavelength spacing, and the dipole numbers of Section VII-B are a strong-coupling stress case for this cell, not a bound on its coupling cost. End to end, the analytic chain with free-space propagation explains a median of 38 to 40% of the configuration-dependent variance per link (31 to 34% with the wall), because it has no element patterns, ground plane or board scattering; that gap is what an estimator trained in simulation has to bridge.

Estimation under hardware impairments. The estimators run on the full-wave channels as in Section III-D: the $M+1$ DFT sweep is applied through the bias table of the calibrated cell, a user moves along a probe track with $T = 10$ observations per window, and the pilot SNR is 10 dB. Four hardware conditions are applied to the simulated observations: none; bias-network parasitics, element tolerances and a 6-bit bias DAC, all unknown to the nominal lookup table, so they act as calibration errors (hw; the via position of integrated bias lines alone moves a 1-bit phase difference by up to 8° at the 29 GHz

design frequency [42]); LO phase noise, IQ imbalance and a 6-bit ADC (rf); and both together (values in Supplementary Section S33). The learned estimators are trained either on an analytic twin of the scene, that is, in simulation, or on full-wave data of the other probe track, that is, on data from the same virtual testbed (Fig. 5(c)). For the simulation-trained estimators the ordering of Section VI-A holds under every condition in both scenes: the deployable Kalman filter leads (0.124 to 0.166 in the clean scene, 0.093 to 0.130 with the wall), the twin-trained LSTM trails it by 16 to 31% (38 to 48% with the wall) and FedCNN by 25 to 42% (37 to 47%), and LS is last (0.216 to 0.258; 0.167 to 0.222), its error nearly double the simulated one because the multiport response departs from the linear cascade that LS assumes. Trained on full-wave data from the same virtual testbed, the LSTM instead comes within about 10% of the Kalman filter in seven of the eight conditions (from 10% below to 4% above its NMSE) and is 18% behind it in the eighth, the wall scene with bias-network impairments. This is the first setting in this study in which a learned estimator comes this close to the classical filter at the corrected interval. Its gain over the twin-trained network, 21 to 32% in the two scenes, is the testbed-specific systematic error, which the classical filters do not model and the analytic twin does not reproduce.

Probe presence. The 84-port solve holds all 48 probes at once, and the channels above are read at one probe with the other 47 open-circuited, so that their port currents are zero, although their conductors still scatter. Eight further solves with a single probe in the scene, four per scene, support this reading: at every position their channel is closer to it (median difference 2 to 14%) than to a reading with the inactive probes terminated in 50 Ω, as in a switched probe array (3 to 41%); remeshing alone changes the solution by up to 1.5%. Terminated, the probes change the channel by a median of 14% (8% with the wall), but the estimator comparison is robust to a change of this size: no NMSE changes by more than 10%, the Kalman filter still leads the simulation-trained estimators with LS last, and the full-wave-trained LSTM again comes within about 10% of it in seven of the eight conditions (Supplementary Section S33).

VIII. LINK-LEVEL AND DEPLOYMENT IMPLICATIONS

A. PILOT OVERHEAD AGAINST THE COHERENCE BUDGET

Converted into time (Supplementary Section S24, Table S10), the overhead is decisive at moderate mobility: with $T_p = 71.35\,\mu$s and Clarke's coherence rule, the full sweep of the $M = 32$ surface fits inside one coherence interval only below 55 km/h ($N_{\mathrm{coh}} = 60$ and $\eta_{\mathrm{net}} = 0.45$ at 30 km/h) and not at 60 or 120 km/h, where the net rate under this protocol is zero. Comparisons at these speeds are model diagnostics rather than deployable operating points, and the overhead argues for element grouping [5] and for compressive designs built on a few active sensing elements [22], independently of the estimator. Spectral efficiencies in Section VIII-B are reported both raw and net of this overhead.

B. LINK-LEVEL IMPACT

Fig. 6 evaluates the estimators inside the precoded QPSK downlink of (2), with the precoder of (22), the pilot SNR fixed at 10 dB and the data SNR swept. Estimation error surfaces as residual multi-user interference, so BER differences reflect the *structure* of the estimation error rather than its magnitude alone. Under ZF the estimator ordering is close to inverted with respect to channel NMSE: at 10 dB data SNR the simulated BERs are 0.0707 (FedCNN), 0.0713 (deployable Kalman), 0.0722 (AR(1)-oracle Kalman), 0.0725 (LS), 0.0876 (selected FL-LSTM), 0.0923 (untuned) and 0.0935 (centralized LSTM), against 0.0254 under perfect CSI; at high data SNR every estimator settles on an error floor (0.031–0.051 at 20 dB) while perfect CSI reaches 5.4×10^{-6}. FedCNN has the worst channel NMSE of any learned method, 25% above the selected estimator, yet the best BER at both 10 and 20 dB. The precoder study at the end of this subsection shows why: most of a learned estimator's NMSE is a gain error along the true channel, to which ZF is insensitive, and FedCNN has the smallest error component orthogonal to the channel, the part that leaks through the ZF nulls. The selected estimator makes the opposite trade and ranks fifth of seven. Together with the field-level result of Section VII-D this indicates that channel NMSE is a misleading figure of merit here, although the field-level comparison is the better-controlled of the two. The primary uncertainty statement is a paired cluster bootstrap on per-run bit counts, obtained by re-running the experiment under the identical protocol and seeds; the intraclass correlation of per-run BER is 0.73 to 0.89, so the 30 runs carry the information of roughly eleven independent ones, and pooled Wilson intervals would understate the uncertainty (Supplementary Section S16). Under this bootstrap the inversion holds within the learned family and the deficit against the Kalman filter holds in full: FedCNN's BER advantage over the selected estimator (-1.7×10^{-2} at 10 dB, -1.5×10^{-2} at 20 dB) and the selected estimator's deficit against the deployable Kalman filter ($+1.6 \times 10^{-2}$ and $+1.3 \times 10^{-2}$) hold with 95% intervals clear of zero. FedCNN's margin over the classical estimators does not: against the deployable Kalman filter the paired difference is -5.6×10^{-4} at 10 dB with interval $[-3.6, +2.5] \times 10^{-3}$ and -2.3×10^{-3} at 20 dB with interval $[-4.6 \times 10^{-3}, +8.8 \times 10^{-5}]$, and against LS it resolves at 20 dB (-4.4×10^{-3}, interval $[-6.7, -2.2] \times 10^{-3}$) but not at 10 dB. "Best BER in the table" is therefore a statement about the learned family; against the classical filters the correct reading is that the best learned estimator is comparable to them at the link level, not consistently better. Sum spectral efficiency and outage under ZF and RZF tell the same story (Supplementary Section S22, Fig. S11): at 10 dB data SNR the ZF sum spectral efficiency spans 7.54 to 8.41 b/s/Hz across estimators against 11.12 under perfect CSI, so the cost of imperfect estimation dominates the spread between estimators, and charging the pilot overhead through (18) moves every method down by more than that spread.

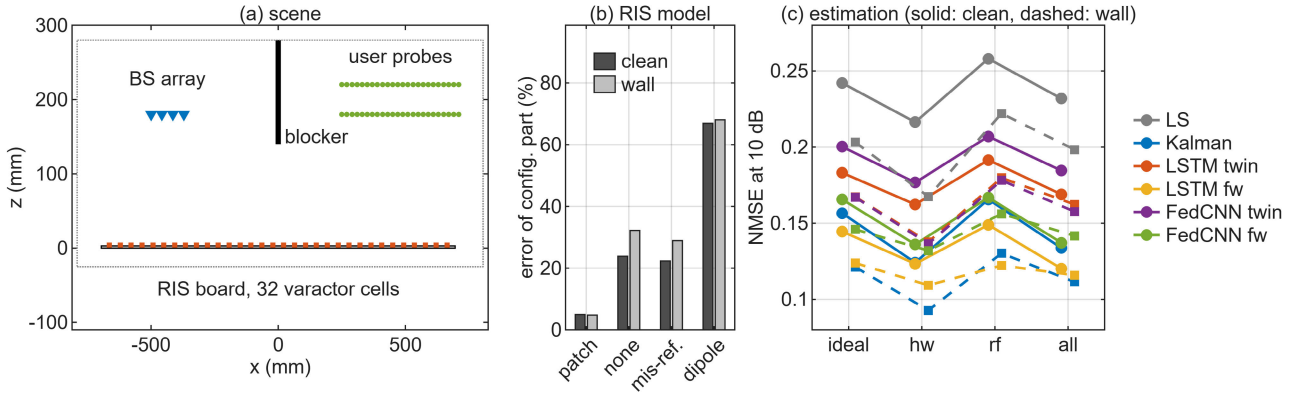


FIGURE 5. Virtual over-the-air testbed. (a) Scene, solved as one 84-port network. (b) Error of the configuration-dependent channel with the RIS model swapped inside the full-wave scene. (c) NMSE at 10 dB under four hardware conditions, inactive probes open-circuited; twin: trained in simulation, fw: trained on full-wave data of the other probe track.

The same mismatch under MRT, ZF and RZF. The effective-channel estimates were re-evaluated at the main-protocol budget (three data seeds, pilot SNR 10 dB) under maximum-ratio transmission (MRT), ZF and RZF, with the error of each estimate split into the part along the true channel and the part orthogonal to it (Supplementary Section S30). The parallel part is a gain error that power normalization cancels under MRT and ZF: scaling LS by its NMSE-optimal scalar lowers NMSE by 8.5% and leaves the MRT and ZF sum spectral efficiency unchanged to four decimals, while the scale-sensitive RZF gains 1.5%. For the learned estimators this invisible part is most of the NMSE, 67% for the selected estimator and 75% for FedCNN, against 7% for LS and 19% for the deployable Kalman filter. FedCNN and the selected estimator both have NMSE 0.214, yet at 10 dB data SNR the selected estimator leads under MRT (6.60 against 6.22 b/s/Hz) and FedCNN under ZF (8.32 against 8.04): FedCNN has the smaller orthogonal error, which ZF leaks, but the larger angle to the channel, which lowers the MRT gain. The lowest-NMSE estimator, the deployable Kalman filter (0.086), is not the best under MRT or ZF either; LS is (7.22 and 8.44 against 7.16 and 8.41 b/s/Hz). Channel NMSE thus misranks estimators under each precoder, and the misranked pairs depend on the precoder through the error structure quantified in Section VII-I.

C. BEAM-SCANNING BASE STATIONS

Beam-scanning leaky-wave antennas steer their beams (two in [43], four in [44]) from a single RF input, by frequency [43], [44] or by one varactor bias shared by all cells [43], and PIN-diode metasurfaces can switch an antenna between omnidirectional and directional patterns [45]. A base station built on them has one RF chain and observes $g = \mathbf{c}^H\mathbf{h}$ for its aperture distribution $\mathbf{c}$. The framework carries over: the LS estimate is linear in the received pilots, so the beam-domain observation is exactly $\mathbf{c}^H\hat{\mathbf{h}}_t$ and every estimator applies with $N = 1$, and the loads that steer the beam are ports of the same network as the RIS loads. Two

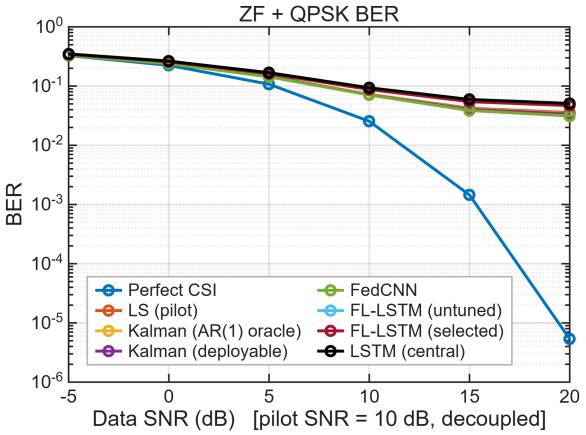


FIGURE 6. Uncoded QPSK BER versus data SNR (ZF precoding; pilot SNR fixed at 10 dB). No operating point was error-free.

tests check this (Supplementary Section S31). In the analytic chain, a leaky-wave aperture steered at the RIS adds 7.2 dB, which moves the operating point to where the retrained learned estimators reach their floor: at 10 dB, LS (0.016) and the Kalman filter (0.017) lead FedCNN (0.059) and the selected estimator (0.064), and the QPSK error rate again follows the phase error rather than the NMSE. A frequency-scanning antenna squints across the band. Scaling the angular dispersion of [43] and [44] from their own bands (5.6 to 6.3 and 7.7 to 10.2 GHz) to 3.5 GHz, a stress test rather than a prediction for those antennas, puts the edge of a 100 MHz carrier 6.1° and 3.1° off the RIS and lowers the beam gain from 7.2 to 0.5 and 5.2 dB. In the full-wave testbed, feeding one base-station dipole and terminating the other three with the RIS varactor under one bias gives a single-feed antenna. Its bias moves the link power by up to 7.9 dB, and the best setting comes within 1.6 dB of four-chain transmission. The parasitic currents must come from the coupled network, since neglecting their mutual coupling misses 73 to 77% of the bias-dependent channel; a 50 Ω load reference costs under 3% because the dipole self-impedance (78 Ω) is close to it. With

this antenna the Kalman filter is the most accurate estimator in five of eight hardware conditions. The LSTM trained on full-wave data leads in the other three, all in the clean scene and by at most 2%, but trails the Kalman filter by 26 to 117% under multipath. A single RF chain gives up the per-antenna channel that ZF needs, so users are served by separate beams [43] or in time.

D. COMMUNICATION EFFICIENCY AND QUANTIZED TRANSMISSION

With traffic measured over the 50-round protocol and per-tensor metadata included (Supplementary Section S6), 8- and 16-bit quantization are indistinguishable from 32 bits in NMSE, while 4 bits costs +9% on the uplink alone and +16% in both directions. The selected model transfers 43.1 MB over 50 rounds at 32/32 bits and 10.8 MB at 8/8, factors of 2.9 and 11.7 below the 32-bit FedCNN benchmark, and federated training undercuts raw-sample upload beyond 30 582 training samples at 32/32 bits and 7 657 at 8/8.

E. COMPUTATIONAL COMPLEXITY

Per inference the LSTM costs $\mathcal{O}(4TH(H+4N))$ operations. The selected model has 26 912 parameters, fewer by factors of 7.5 and 2.9 than the centralized CNN-LSTM and FedCNN, an inference latency of 1.089 ms per sample against 5.418 and 2.584 ms, and a 50-round federated training time of 495 ± 85 s; the one-time search and playoff took 1 263 CPU-minutes. An LMMSE estimate of the cascaded channel would instead invert an $N(M+1) = 528$-dimensional covariance, at $\mathcal{O}(N^3(M+1)^3)$ operations [7]. The measured table is in Supplementary Section S12.

IX. SENSITIVITY ANALYSES AND LIMITATIONS

Seven limitations bound the claims.

(i) *No fabricated hardware.* The evidence comes from the analytic chain and from full-wave solves. The virtual testbed of Section VII-J tests whether the conclusions persist beyond the analytic propagation model, but it uses ideal lumped ports instead of packaged diodes and a bias board (lumped-element diode models moved the reflection band of a 29 GHz cell by up to 10 GHz against a three-dimensional diode model [41], so a fabricated board needs the calibration of Section VII-E repeated on measured states), a scene truncated at ±55 mm transverse to the array, two probe tracks and a single frequency; a measured channel remains future work. The hardware sweeps of the analytic chain (phase quantization, incidence and polarization, dispersion, component tolerances, a patch base-station array) are reported over four seeds in Supplementary Section S32 and are not folded into the main comparisons. They leave the ordering of the main comparison intact in every condition except 1-bit phase resolution, where the learned estimator is the most accurate in three of four repetitions.

(ii) *One cell topology.* The five cell families of Section VII-H vary loss, tuning range and circuit, but share the two-patch topology on one substrate, and the patch kernel describes that cell. The bias table is calibrated at normal incidence, and the calibrated circuit departs from HFSS beyond 45° incidence. Cells that switch between reflection and transmission [41], [42] lie outside this reflective model.

(iii) *The pilot protocol is the worst case.* We sweep all $M+1$ patterns per observation. Element grouping [5] and compressive designs with a few active sensing elements [22] shorten or replace the sweep, so the feasibility statements apply to the full sweep. Because T_{obs} grows with M, the M sweeps vary aperture and decorrelation together, and the window $T = 10$ sits past the optimum of $T = 5$ (Section VI-C).

(iv) *Labels and data locality.* Teacher labels are modeled, not operationally validated. Data locality refers to raw observations rather than to a cryptographic guarantee: FedAvg exchanges plaintext updates that can leak information [11], and differential privacy [20] or secure aggregation would cost accuracy or efficiency [11] that we did not measure. Site heterogeneity is statistical [29] and excludes stragglers and asynchrony [11].

(v) *Estimator coverage.* No compared estimator models the coupling kernel; a kernel-aware LS or the formulations of [24], [25], [26] would recover part of the coupling cost. Vector-state and spatio-temporal classical filters were not tried, and neither were networks conditioned on the pilot SNR; the learned estimators are trained once across the SNR grid, which shapes their high-SNR floor. On the testbed, an LSTM trained on full-wave data from that testbed comes within about 10% of the Kalman filter in seven of eight conditions, so the ranking we report is a ranking of estimators trained in simulation.

(vi) *Statistical scope.* The field-level tests are scene-paired and quoted from one training seed, with their sign checked over three. The link-level comparison resolves within the learned family and against the Kalman filter, but not FedCNN against the classical filters (Section VIII-B). The mechanism, adaptation and testbed studies use one to three training seeds.

(vii) *System scope.* All results are confined to our single-carrier MU-MISO downlink; the NOMA and massive MIMO settings of [9] and [10], multi-antenna users and wideband operation remain open.

X. CONCLUSION

This paper examined the consequences of idealizations in the evaluation of RIS channel estimators by rebuilding the pipeline around the physics and the pilot protocol rather than around the estimator. Within that pipeline the federated LSTM is communication-efficient and reproducible to configure, but it is not the most accurate estimator: a Kalman filter that estimates its own parameters from the observations it receives leads it wherever consecutive observations are not strongly correlated. The study yields four findings.

First, referring the cell reflection to the wave impedance inside the multiple-scattering inverse models a passive surface as active without any numerical symptom, and it inflates the error of model-inverting estimators far more than that of the learned one in all five cell families and under both kernels.

The full-wave testbed confirms that the self-referenced form is the exact network-theoretic reduction of its multiport model and that the mis-referenced one is not. Passivity should be verified rather than assumed, at the cost of one sign test on the implied element loads.

Second, spacing observations by a nominal slot instead of the full sweep overstates temporal correlation. Retrained for the corrected interval, the recurrent estimator has 57% higher NMSE than its nominal-slot counterpart, against 5% for the deployable Kalman filter. The ranking follows the normalized Doppler, and correcting the protocol moves the 30 km/h operating point of an $M = 32$ surface into the regime where the classical filter leads.

Third, channel NMSE misranks estimators. Despite a nearly tenfold lead over LS in column NMSE at half-wavelength spacing in the field study, and a 3.6- to 6.6-fold lead in all 60 combinations of configuration, beamforming or RIS-control strategy and training seed, the learned estimator delivers less field at the user, because its residual survives the matched-filter projection and is mostly amplitude error that a phase-only surface ignores. The error-phase coherence, an offline diagnostic like NMSE, predicts the delivered field, and at link level the error orthogonal to the channel, rather than NMSE, largely determines performance under MRT, ZF and RZF. Estimators used for RIS configuration should be evaluated on the delivered field.

Fourth, on a full-wave testbed with bias-network parasitics, tolerances, calibration and RF impairments and multipath, estimators trained in simulation keep the simulated ranking, while an LSTM trained on full-wave data from the same virtual testbed comes within about 10% of the Kalman filter in seven of eight conditions. Its advantage over the twin-trained LSTM comes from the testbed-specific systematic error, which testbed-matched full-wave data contain and the analytic twin does not.

On the method, the origin-blind playoff selected a hand-tuned configuration over the Pareto search because the five-round proxy ranked the eventual winner 162nd of 228; a staged-fidelity search in the manner of successive halving within Hyperband [34], with federated rounds as the allocated resource, is a natural way to improve the search. The next steps follow from Section IX: measurement on fabricated hardware, reduced-overhead pilots evaluated against the coherence bound, kernel-aware and spatio-temporal classical estimators, and training objectives aligned with the configuration task rather than with channel NMSE.

APPENDIX A
STANDARD BUILDING BLOCKS

The equations below are standard and collected here for completeness; the main text refers to them where used.

LSTM cell with forget gate [35], [36]. With input $\mathbf{x}_t$ and $\mathbf{z}_t = [\mathbf{s}_{t-1}; \mathbf{x}_t]$,

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{z}_t + \mathbf{b}_f), \quad \mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{z}_t + \mathbf{b}_i),$$
$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{z}_t + \mathbf{b}_o), \tag{25}$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_c \mathbf{z}_t + \mathbf{b}_c), \tag{26}$$
$$\mathbf{s}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t). \tag{27}$$

Weighted federated averaging [12]. With $\mathcal{S}_r$ the sites participating in round r and n_c the sample count of site c,

$$\mathbf{w}^{(r+1)} = \sum_{c \in \mathcal{S}_r} \frac{n_c}{\sum_{j \in \mathcal{S}_r} n_j} \mathbf{w}_c^{(r)}. \tag{28}$$

The weights are normalized over the participating sites, as in the corrected pseudocode of [12] (arXiv:1602.05629v4); the version printed in the proceedings sums over all sites, which agrees with (28) only when every site participates.

Particle swarm update [31], [32], in the inertia-weight form of [32]. With $r_1, r_2 \sim U(0, 1)$, personal best $\mathbf{p}_{\text{best},p}$ and leader $\boldsymbol{\ell}_p$ (an archive member as in [37], in place of the global best of [31], [32]),

$$\mathbf{v}_p \leftarrow w\,\mathbf{v}_p + c_1 r_1(\mathbf{p}_{\text{best},p} - \mathbf{x}_p) + c_2 r_2(\boldsymbol{\ell}_p - \mathbf{x}_p), \tag{29}$$
$$\mathbf{x}_p \leftarrow \mathbf{x}_p + \mathbf{v}_p. \tag{30}$$

ACKNOWLEDGMENT

During the preparation of this paper, the authors used ChatGPT (OpenAI) and Claude (Anthropic) for language refinement throughout the paper and for consistency review. ChatGPT was additionally used as an assistive tool for bibliographic-metadata checking and for suggesting consistency and validation checks. All simulations, full-wave computations, numerical analyses, verification against primary sources, and final scientific interpretations were performed or independently verified by the authors. The authors reviewed all AI-assisted material and take full responsibility for the paper.

REFERENCES

[1] Q. Wu and R. Zhang, "Intelligent reflecting surface enhanced wireless network via joint active and passive beamforming," *IEEE Trans. Wireless Commun.*, vol. 18, no. 11, pp. 5394–5409, Nov. 2019, doi: 10.1109/TWC.2019.2936025.

[2] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M.-S. Alouini, and R. Zhang, "Wireless communications through reconfigurable intelligent surfaces," *IEEE Access*, vol. 7, pp. 116753–116773, 2019, doi: 10.1109/ACCESS.2019.2935192.

[3] M. Di Renzo, A. Zappone, M. Debbah, M.-S. Alouini, C. Yuen, J. de Rosny, and S. Tretyakov, "Smart radio environments empowered by reconfigurable intelligent surfaces: How it works, state of research, and the road ahead," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 11, pp. 2450–2525, Nov. 2020, doi: 10.1109/JSAC.2020.3007211.

[4] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah, and C. Yuen, "Reconfigurable intelligent surfaces for energy efficiency in wireless communication," *IEEE Trans. Wireless Commun.*, vol. 18, no. 8, pp. 4157–4170, Aug. 2019, doi: 10.1109/TWC.2019.2922609.

[5] B. Zheng and R. Zhang, "Intelligent reflecting surface-enhanced OFDM: Channel estimation and reflection optimization," *IEEE Wireless Commun. Lett.*, vol. 9, no. 4, pp. 518–522, Apr. 2020, doi: 10.1109/LWC.2019.2961357.

[6] Z. Wang, L. Liu, and S. Cui, "Channel estimation for intelligent reflecting surface assisted multiuser communications: Framework, algorithms, and analysis," *IEEE Trans. Wireless Commun.*, vol. 19, no. 10, pp. 6607–6620, Oct. 2020, doi: 10.1109/TWC.2020.3004330.

[7] N. K. Kundu and M. R. McKay, "Channel estimation for reconfigurable intelligent surface aided MISO communications: From LMMSE to deep learning solutions," *IEEE Open J. Commun. Soc.*, vol. 2, pp. 471–487, 2021, doi: 10.1109/OJCOMS.2021.3063171.

[8] H. Ye, G. Y. Li, and B.-H. Juang, "Power of deep learning for channel estimation and signal detection in OFDM systems," *IEEE Wireless Commun. Lett.*, vol. 7, no. 1, pp. 114–117, Feb. 2018, doi: 10.1109/LWC.2017.2757490.

[9] C. Nguyen, T. M. Hoang, and A. A. Cheema, "Channel estimation using CNN-LSTM in RIS-NOMA assisted 6G network," *IEEE Trans. Mach. Learn. Commun. Netw.*, vol. 1, pp. 43–60, 2023, doi: 10.1109/TMLCN.2023.3278232.

[10] A. M. Elbir and S. Coleri, "Federated learning for channel estimation in conventional and RIS-assisted massive MIMO," *IEEE Trans. Wireless Commun.*, vol. 21, no. 6, pp. 4255–4268, Jun. 2022, doi: 10.1109/TWC.2021.3128392.

[11] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, May 2020, doi: 10.1109/MSP.2020.2975749.

[12] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, vol. 54, Fort Lauderdale, FL, USA, 2017, pp. 1273–1282. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[13] W. Ni, Y. Liu, Z. Yang, H. Tian, and X. Shen, "Federated learning in multi-RIS-aided systems," *IEEE Internet Things J.*, vol. 9, no. 12, pp. 9608–9624, Jun. 15, 2022, doi: 10.1109/JIOT.2021.3130444.

[14] A. M. Elbir, A. K. Papazafeiropoulos, and S. Chatzinotas, "Federated learning for physical layer design," *IEEE Commun. Mag.*, vol. 59, no. 11, pp. 81–87, Nov. 2021, doi: 10.1109/MCOM.101.2100138.

[15] G. Gradoni and M. Di Renzo, "End-to-end mutual coupling aware communication model for reconfigurable intelligent surfaces: An electromagnetic-compliant approach based on mutual impedances," *IEEE Wireless Commun. Lett.*, vol. 10, no. 5, pp. 938–942, May 2021, doi: 10.1109/LWC.2021.3050826.

[16] M. Di Renzo, F. H. Danufane, and S. Tretyakov, "Communication models for reconfigurable intelligent surfaces: From surface electromagnetics to wireless networks optimization," *Proc. IEEE*, vol. 110, no. 9, pp. 1164–1209, Sep. 2022, doi: 10.1109/JPROC.2022.3195536.

[17] X. Qian and M. Di Renzo, "Mutual coupling and unit cell aware optimization for reconfigurable intelligent surfaces," *IEEE Wireless Commun. Lett.*, vol. 10, no. 6, pp. 1183–1187, Jun. 2021, doi: 10.1109/LWC.2021.3061449.

[18] P. Zheng, X. Ma, and T. Y. Al-Naffouri, "On the impact of mutual coupling on RIS-assisted channel estimation," *IEEE Wireless Commun. Lett.*, vol. 13, no. 5, pp. 1275–1279, May 2024, doi: 10.1109/LWC.2024.3367913.

[19] R. Faqiri, C. Saigre-Tardif, G. C. Alexandropoulos, N. Shlezinger, M. F. Imani, and P. del Hougne, "PhysFad: Physics-based end-to-end channel modeling of RIS-parametrized environments with adjustable fading," *IEEE Trans. Wireless Commun.*, vol. 22, no. 1, pp. 580–595, Jan. 2023, doi: 10.1109/TWC.2022.3196834.

[20] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Vienna, Austria, Oct. 2016, pp. 308–318, doi: 10.1145/2976749.2978318.

[21] C.-J. Chun, J.-M. Kang, and I.-M. Kim, "Deep learning-based channel estimation for massive MIMO systems," *IEEE Wireless Commun. Lett.*, vol. 8, no. 4, pp. 1228–1231, Aug. 2019, doi: 10.1109/LWC.2019.2912378.

[22] A. Taha, M. Alrabeiah, and A. Alkhateeb, "Enabling large intelligent surfaces with compressive sensing and deep learning," *IEEE Access*, vol. 9, pp. 44304–44321, 2021, doi: 10.1109/ACCESS.2021.3064073.

[23] M. Di Renzo and P. del Hougne, "Multiport network theory for modeling and optimizing reconfigurable metasurfaces," in *Proc. 19th Eur. Conf. Antennas Propag. (EuCAP)*, Stockholm, Sweden, Mar. 2025, pp. 1–4, doi: 10.23919/eucap63536.2025.10999590.

[24] P. Zheng, S. Tarboush, H. Sarieddeen, and T. Y. Al-Naffouri, "Mutual coupling-aware channel estimation and beamforming for RIS-assisted communications," *IEEE Trans. Wireless Commun.*, vol. 25, pp. 4374–4390, 2026, doi: 10.1109/TWC.2025.3610587.

[25] T. Qiu, R. Li, C. Pan, T. Zhang, D. Xia, C. Wang, and H. Ren, "Mutual coupling aware channel estimation for RIS-aided multi-user mmWave systems," *IEEE Trans. Wireless Commun.*, vol. 25, pp. 15803–15817, 2026, doi: 10.1109/TWC.2026.3686033.

[26] A. Dkhan, S. Tarboush, H. Sarieddeen, and T. Y. Al-Naffouri, "RIS-aided near-field channel estimation under mutual coupling and spatial correlation," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Taipei, Taiwan, Dec. 2025, pp. 599–604, doi: 10.1109/GLOBECOM59602.2025.11432750.

[27] M. Nerini, H. Li, and B. Clerckx, "Global optimal closed-form solutions for intelligent surfaces with mutual coupling: Is mutual coupling detrimental or beneficial?" *IEEE Trans. Wireless Commun.*, vol. 25, pp. 3201–3214, 2026, doi: 10.1109/TWC.2025.3602218.

[28] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," 2016, *arXiv:1610.05492*.

[29] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-IID data," 2018, *arXiv:1806.00582*.

[30] Z. Li, S. Li, U. Awada, and J. Zhang, "Joint channel estimation and dynamics-aware grouping for time-varying RIS-assisted OTA federated learning," 2026, *arXiv:2607.17169*.

[31] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. Int. Conf. Neural Netw.*, vol. 4, Perth, WA, Australia, 1995, pp. 1942–1948, doi: 10.1109/icnn.1995.488968.

[32] Y. Shi and R. Eberhart, "A modified particle swarm optimizer," in *Proc. IEEE Int. Conf. Evol. Comput. (IEEE World Congr. Comput. Intell.)*, Anchorage, AK, USA, May 1998, pp. 69–73, pp. 69–73, doi: 10.1109/ICEC.1998.699146.

[33] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, no. 10, pp. 281–305, 2012. [Online]. Available: https://jmlr.org/papers/v13/bergstra12a.html

[34] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *J. Mach. Learn. Res.*, vol. 18, no. 185, pp. 1–52, 2018. [Online]. Available: https://jmlr.org/papers/v18/16-558.html

[35] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[36] F. A. Gers, J. Schmidhuber, and F. Cummins, "Learning to forget: Continual prediction with LSTM," *Neural Comput.*, vol. 12, no. 10, pp. 2451–2471, Oct. 2000, doi: 10.1162/089976600300015015.

[37] C. A. C. Coello, G. T. Pulido, and M. S. Lechuga, "Handling multiple objectives with particle swarm optimization," *IEEE Trans. Evol. Comput.*, vol. 8, no. 3, pp. 256–279, Jun. 2004, doi: 10.1109/TEVC.2004.826067.

[38] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002, doi: 10.1109/4235.996017.

[39] Q. H. Spencer, A. L. Swindlehurst, and M. Haardt, "Zero-forcing methods for downlink spatial multiplexing in multiuser MIMO channels," *IEEE Trans. Signal Process.*, vol. 52, no. 2, pp. 461–471, Feb. 2004, doi: 10.1109/TSP.2003.821107.

[40] C. B. Peel, B. M. Hochwald, and A. L. Swindlehurst, "A vector-perturbation technique for near-capacity multiantenna multiuser communication—Part I: Channel inversion and regularization," *IEEE Trans. Commun.*, vol. 53, no. 1, pp. 195–202, Jan. 2005, doi: 10.1109/TCOMM.2004.840638.

[41] I. Gulen, V. R. Kandregula, S. Chalkidis, S. Rizou, P. Lazaridis, A. Feresidis, and E. Vassos, "A dual function metasurface based on aluminium gallium arsenide PIN diodes," in *Proc. 28th Int. Symp. Wireless Pers. Multimedia Commun. (WPMC)*, Sofia, Bulgaria, Nov. 2025, pp. 1–6, doi: 10.1109/WPMC67460.2025.11351032.

[42] I. Gulen, E. Vassos, V. R. Kandregula, S. Chalkidis, S. Rizou, P. Lazaridis, and A. Feresidis, "A dual function metasurface for RIS applications," in *Proc. 20th Eur. Conf. Antennas Propag. (EuCAP)*, Dublin, Ireland, Apr. 2026, pp. 1–5, doi: 10.23919/EUCAP68105.2026.11611932.

[43] S. Wang, W. Wang, and Y. Zheng, "Dual-functional quasi-uniform beam-scanning antenna array with endfire radiation capability for integrated sensing and communication applications," *IEEE Trans. Veh. Technol.*, vol. 74, no. 11, pp. 17829–17839, Nov. 2025, doi: 10.1109/TVT.2025.3580083.

[44] S. Wang, W. Wang, K. L. Chung, and Y. Zheng, "An orthogonal quad-beam scanning antenna using 1-bit dielectric modulation in plasmonic metamaterial transmission line for traffic monitoring applications," *IEEE Trans. Veh. Technol.*, vol. 75, no. 2, pp. 2159–2168, Feb. 2026, doi: 10.1109/TVT.2025.3598024.

[45] V. R. Kandregula, P. I. Lazaridis, I. Gulen, S. Chalkidis, A. Feresidis, V. Poulkov, Z. D. Zaharis, and E. Vassos, "Design process for a dual pattern reconfigurable antenna based on flip-chip aluminum gallium arsenide (AlGaAs) PIN diodes," in *Proc. 28th Int. Symp. Wireless Pers. Multimedia Commun. (WPMC)*, Sofia, Bulgaria, Nov. 2025, pp. 1–6, doi: 10.1109/WPMC67460.2025.11351125.



NURI GÖKMEN KARAKIRAZ received the B.Sc. degree in electrical and electronics engineering from the Faculty of Engineering, Ankara University, Ankara, Türkiye, in 1997, the M.Sc. degree in electrical and electronics engineering from Hacettepe University, Ankara, in 2010, with a thesis on null steering in array antennas using genetic algorithms, and the Ph.D. degree in electronics engineering from National Defense University, Türkiye, in 2022, with research on microwave imaging applications. He began his academic career as a Lecturer with the Air Force Academy, National Defense University, where he worked as the Department Head and the Division Head. He is currently an Assistant Professor with the Department of Electrical and Electronics Engineering, Faculty of Engineering, Istanbul Aydin University, Istanbul, Türkiye. He has authored articles published in international refereed journals. His research interests include microwave imaging, electromagnetic inverse problems, array antennas, and optimization methods.



AGAH OKTAY ERTAY (Member, IEEE) received the B.Sc. degree in electronics and communication engineering from Yıldız Technical University, Istanbul, Türkiye, in 2010, and the M.Sc. and Ph.D. degrees in telecommunications engineering from Istanbul Technical University (ITU), Istanbul, in 2014 and 2019, respectively.

He was a Research Assistant with ITU, from 2011 to 2019. Since 2019, he has been an Assistant Professor with the Department of Electrical and Electronics Engineering, Erzincan Binali Yıldırım University. His research interests include RF and microwave engineering, electromagnetic wave propagation, antennas and propagation, passive microwave circuits, periodic and electromagnetic bandgap structures, slow-wave structures, and photonic sensing.

Dr. Ertay is a member of the IEEE Antennas and Propagation Society. He received the IEEE Ulrich L. Rohde Innovative Conference Paper Award at the 2016 IEEE Conference on Antenna Measurements and Applications (CAMA). He is a Reviewer of IEEE Access, IEEE Transactions on Industrial Electronics, IEEE Transactions on Electron Devices, IEEE Transactions on Microwave Theory and Techniques, IEEE Microwave and Wireless Components Letters, IEEE Sensors Letters, IEEE Open Journal of Antennas and Propagation, and IEEE Antennas and Wireless Propagation Letters.



UMUT ŞÜKRÜ YAŞAR received the B.Sc. degree in electrical and electronics engineering from Eskişehir Osmangazi University, Eskişehir, Türkiye, in 2013, the M.Sc. degree in electrical and electronics engineering from Erzincan Binali Yıldırım University, Erzincan, Türkiye, in 2018, and the Ph.D. degree in electrical and electronics engineering from Karadeniz Technical University, Trabzon, Türkiye, in 2025. His doctoral research focused on electromagnetic shielding. He is currently an Assistant Professor with the Department of Motor Vehicles and Transportation Technologies, Kemah Vocational School, Erzincan Binali Yıldırım University. His research interests include electromagnetic fields and microwave engineering.

HALIT NAZLI received the B.Sc. and M.Sc. degrees in electrical and electronics engineering from Istanbul Aydin University, Istanbul, Türkiye, in 2020 and 2022, respectively, and the Ph.D. degree in electrical and electronics engineering from Istanbul Aydin University, in 2026. His research interests include fuzzy logic systems, image processing, skeletonization, shape analysis, medical image analysis, and intelligent decision-making methods.

SAMAR A. S. AL-HASASNEH received the B.Sc. degree in computer science from the University of Sharjah, United Arab Emirates, and the M.Sc. degree in information systems from Ajman University of Science and Technology, United Arab Emirates, in 2013. She is currently a Computer Lecturer with Al Wasl University, Dubai, United Arab Emirates. Her research interests include data warehousing, databases, data mining, software engineering, global positioning systems (GPS) tracking systems, and radio frequency identification (RFID) technology.