JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

HYBRID NEAR-FIELD AND FAR-FIELD CHANNEL ESTIMATION FOR RIS-ASSISTED mmWAVE MASSIVE MIMO SYSTEMS USING GRU WITH ATTENTION

Satya Chinmayi M, Sailaja M
Department of Electronics and Communication Engineering
University College of Engineering, Kakinada
Andhra Pradesh, India

Abstract : *5G and 6G networks are being built around three technologies that rarely appeared together in earlier standards: millimetre-wave carriers, massive antenna arrays, and reconfigurable intelligent surfaces (RIS). Put these together and channel estimation stops being routine. Users and scatterers no longer sit cleanly in the far field; many fall into a near-field region where the wavefront curves, and a single link can carry both kinds of propagation at once. This paper works through that problem for an RIS-assisted mmWave massive MIMO system, treating the near- and far-field contributions as one hybrid channel instead of two separate estimation problems. The channel is built using OFDM pilot transmission and RIS-controlled reflection, and 5000 labelled samples are generated from this pipeline to train three different estimators on them - a K-nearest-neighbour (KNN) classifier as a simple baseline, a gated recurrent unit (GRU) network, and a GRU model with an added attention layer. The GRU picks up the sequential structure in the channel data; attention then tells the network which parts of that structure actually matter before it makes a prediction. Together, these form a practical framework for exploring learning-based channel estimation in the kind of mixed-propagation environment that RIS-assisted 5G/6G networks are expected to operate in.*

Index Terms *- Millimeter wave (mmWave), massive MIMO, reconfigurable intelligent surface (RIS), near-field propagation, far-field propagation, hybrid channel estimation, OFDM, KNN, GRU, attention mechanism.*

I. INTRODUCTION

6G research keeps returning to one idea: make the antenna array much bigger. Bigger arrays promise higher throughput and better spectral efficiency, which is why extremely large-scale MIMO has become such an active research area. But bigger arrays change the physics of the problem. As the number of antennas climbs into the thousands, the Rayleigh distance - the point at which propagation stops looking planar and starts looking spherical - stretches out to hundreds of metres. A large share of users and scatterers end up inside this near-field zone rather than beyond it, and once that happens, the usual assumption of flat, plane-wave propagation no longer holds. Channel estimation then has to recover angle and distance together, not one after the other, and that joint search is far more expensive computationally than anything a conventional far-field method or a standard subspace search was built to handle, especially once the array is planar rather than linear [3]. Since accurate channel state information (CSI) is what makes the gains from a large array usable in the first place, this difficulty has pushed a lot of recent work toward near-field and hybrid-field estimation - and a fair amount of it is directly relevant here.

Some of that work focuses purely on getting the near-field model right for planar arrays. One line of research builds an end-to-end spherical-wavefront model that represents non-line-of-sight paths as virtual sources, then pairs it with a low-complexity initialisation routine and a refinement step whose accuracy is backed up theoretically [1]. Another paper takes a different route: rather than estimating angle and distance jointly, it decouples them completely - angles come from a two-dimensional Fourier-based estimator, distance follows from a closed-form expression - and this separation turns out to cut complexity while still matching or beating existing near-field methods on accuracy [3].

Things get harder once scatterers are spread across both distance ranges at once, which is what people usually mean by hybrid-field propagation. Compressive-sensing methods built around a fixed dictionary don't cope well with this: the dictionary atoms are fixed in advance and simply can't bend to fit a scatterer distribution that mixes near and far field, and near-field dictionaries in particular need such dense angle-distance sampling that the true scatterer position is rarely close to any grid point - accuracy drops, computation climbs. One way around this is to learn the dictionary directly from the measured channel data instead of fixing it in advance, solving for both the dictionary and the sparse representation jointly through an ADMM procedure; this consistently beats fixed-dictionary baselines whether the channel is purely near-field or genuinely hybrid [2]. A related paper points out that classical tools like least squares or orthogonal matching pursuit simply aren't built to exploit the coupling between near- and far-field structure, and proposes instead a joint sparse Bayesian learning approach that fuses both dictionaries under one shared prior - this removes the need to know the number of paths ahead of time and holds up better when noise gets difficult [4].

Pilot overhead is its own separate headache as arrays scale up, since the conventional rule of thumb is one pilot per transmit antenna, and that number gets large fast. One way to soften this is to transmit pilots on only a subset of the antennas, estimate that partial channel with a least-mean-square approach, and then let a neural network fill in the rest by exploiting the spatial correlation that naturally exists across the array - the estimation error rises slightly, but throughput improves and bit-error-rate stays roughly where it was [5]. RIS-assisted systems make this worse before they make it better, because the combined channel dimensionality is so much higher; one proposed fix pairs a convolutional network that picks out the most useful RIS elements through a differentiable Top-N operation with an attention-and-residual estimator that then reconstructs the full time-varying channel, and this holds up reasonably well even when users are moving quickly [9].

Some RIS-assisted designs skip the idea of separate estimation stages altogether. One approach treats the combined direct-plus-reflected signal as a third-order tensor and recovers every channel component in one training pass using alternating least squares, without ever needing to switch the RIS off during estimation - pilot cost goes down and the result is noticeably more robust to noise [7]. Holographic MIMO raises a related but distinct problem: apertures are packed so densely that ordinary angular-domain estimation stops working once there is more than one user, so one paper instead moves the whole problem into the wavenumber domain and applies a joint graph-cut algorithm that alternates between finding supports shared across users and refining each user's own estimate - accuracy improves and the pilot sequence gets shorter [8]. A rather different question is asked elsewhere: does jointly optimising estimation and beamforming actually buy anything over doing them separately? Across essentially every standard massive MIMO model, the answer turns out to be no - splitting the problem into MMSE estimation followed by MMSE beamforming loses nothing relative to a joint design, centralised or distributed, which suggests that AI-based methods fusing the two stages together may not have much room to add value once these assumptions hold [6].

Put together, this literature keeps returning to two points: near-field and hybrid-field propagation genuinely change what channel estimation needs to do, and most of the fixes proposed so far still trade complexity for accuracy in ways that don't sit comfortably with RIS-assisted deployments specifically. That gap is what this paper tries to address. Rather than estimating near-field and far-field contributions through separate stages or a hand-built dictionary, near- and far-field contributions are here treated as parts of one hybrid channel, letting a data-driven pipeline learn that structure directly.

The rest of the paper follows this order: Section II sets up the problem formally, Section III walks through the three estimators - KNN, GRU, and GRU with attention - Section IV reports simulation results, and Section V closes with conclusions.

II. PROBLEM STATEMENT

Reconfigurable intelligent surfaces (RIS) are attractive largely because they are cheap relative to what they offer: a second, controllable propagation path between base station and user, sitting alongside whatever direct path already exists. That second path is only useful, though, if the resulting channel can actually be estimated accurately, and that is harder than it sounds once near-field and far-field conditions can both be present in the same link. Near-field propagation follows a spherical wavefront - distance to the scatterer directly affects both the phase and the amplitude seen at each antenna - while far-field propagation is close enough to planar that angle alone describes it reasonably well. An estimator built around only one of these assumptions will misrepresent the channel whenever the other one actually applies.

For the system considered here, the overall channel is the sum of the direct base-station-to-user path and the path that bounces off the RIS:

$$Hhybrid = HBU + HBR\ \Theta\ HRU \tag{1}$$

H_BU is the direct channel, H_BR the base-station-to-RIS channel, H_RU the RIS-to-user channel, and Θ the RIS's phase-shift matrix. With the dimensions used in this work, H_BR is 16 × 64, H_RU is 64 × 1, and H_hybrid ends up 16 × 1.

Θ itself is diagonal - each RIS element only shifts phase, it does not mix signals between elements - so

$$\Theta = diag(ej\theta 1,\ ej\theta 2,\ \dots ,\ ej\theta N) \tag{2}$$

and with N = 64 here, that makes Θ a 64 × 64 diagonal matrix.

Least squares and MMSE are the obvious first choices for estimating H_hybrid, but both start to struggle once the channel dimension grows large and the underlying propagation is a nonlinear mix of near- and far-field behaviour - neither method was designed with that kind of structure in mind. That is the motivation for looking at data-driven alternatives instead. The approach here starts with KNN as a straightforward machine-learning baseline, moves to a GRU network that can pick up temporal and structural dependencies across channel observations, and finishes with a GRU-plus-attention variant that lets the model weight the most informative features more heavily during estimation.

III. PROPOSED ALGORITHM

With the problem set up in Section II, three estimators for the hybrid RIS-assisted channel are considered next, each more capable than the last: K-nearest neighbours (KNN), a gated recurrent unit (GRU), and a GRU extended with attention. All three are trained and tested on the same dataset - generated from an OFDM transmission-and-reception pipeline - so the comparison between a conventional baseline and two deep-learning variants sits on equal footing.

A. K-Nearest Neighbours (KNN)

KNN is the simplest of the three and serves as the baseline. Given a new sample, it looks at the labels of the closest samples in the training set and votes among them. The features here come from the received channel observations. For a test sample x and training sample xi, distance is measured the usual way, via Euclidean distance:

$$d(x, x_i) = \sqrt{\sum_{j=1}^{n} (x_j - x_{ij})^2} \tag{3}$$

with n the number of features. The K closest training points are found, and whichever class shows up most often among them becomes the prediction. In this work that means sorting each channel observation into Low, Medium, or High, purely as a point of comparison for the two deep-learning models that follow.

B. Gated Recurrent Unit (GRU)

A GRU, unlike KNN, is built to handle sequences - which matters here because successive channel observations are not independent of each other. Each GRU cell has two gates doing the work: an update gate that decides how much of the earlier hidden state to keep, and a reset gate that decides how much of it to drop when forming the new candidate state.

The update gate:

$$z_t = \sigma(W_z x_t + U_z h_{t-1}) \tag{4}$$

the reset gate:

$$r_t = \sigma(W_r x_t + U_r h_{t-1}) \tag{5}$$

and from the reset gate, the candidate hidden state:

$$\tilde{h}_t = \tanh(W_h x_t + U_h (r_t \odot h_{t-1})) \tag{6}$$

the update gate then blends this candidate with the previous state to give the final hidden state at time t:

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \tag{7}$$

The GRU here takes the preprocessed OFDM/channel observations one step at a time, picks up whatever temporal pattern is present, and hands the result to a dense output layer that produces the predicted channel state.

C. GRU with Attention Mechanism

The last and most capable model adds attention on top of the GRU so the network is not forced to rely only on the final hidden state - it can weight earlier time steps too, wherever they turn out to matter. Starting from the full sequence of hidden states the GRU produces,

$$H = [h_1, h_2, \dots , h_T] \tag{8}$$

attention assigns each one a score:

$$e_t = v^T \tanh(W h_t + b) \tag{9}$$

normalised across the sequence with softmax:

$$\alpha_t = \exp(e_t) / \sum_{k=1}^{T} \exp(e_k) \tag{10}$$

and the resulting weights combine into a single context vector, summarising the sequence according to how much each step actually mattered:

$$c = \sum_{t=1}^{T} \alpha_t h_t \tag{11}$$

The division of labour is fairly intuitive: the GRU extracts the sequential structure, attention then figures out which part of that structure carries the most useful information for the prediction and weights it up accordingly. The result handles mixed near-field/far-field characteristics noticeably better than a GRU left to rely on just its last hidden state.

IV. RESULTS AND DISCUSSION

Here the three estimators - KNN, GRU, GRU with attention - are compared on the hybrid RIS-assisted channel from Section II, using two measures: NMSE across a spread of SNR values, and confusion matrices for each model's channel-state classification.

A. NMSE Performance

NMSE was computed at several SNR points for each model:

$$NMSE = ||H - \hat{H}||_2^2 \,/\, ||H||_2^2 \tag{12}$$

with H the true channel and Ĥ its estimate. Fig. 1 shows NMSE against SNR for all three. Error falls as SNR rises across the board - unsurprising, since noise matters less once the signal is stronger. KNN sits highest throughout, which tracks with its limited ability to model channel structure at all. GRU does noticeably better, since it is actually learning temporal patterns from the data rather than just measuring distance. GRU with attention comes out lowest, which fits the idea that focusing on the more informative parts of the sequence gives it a real edge.

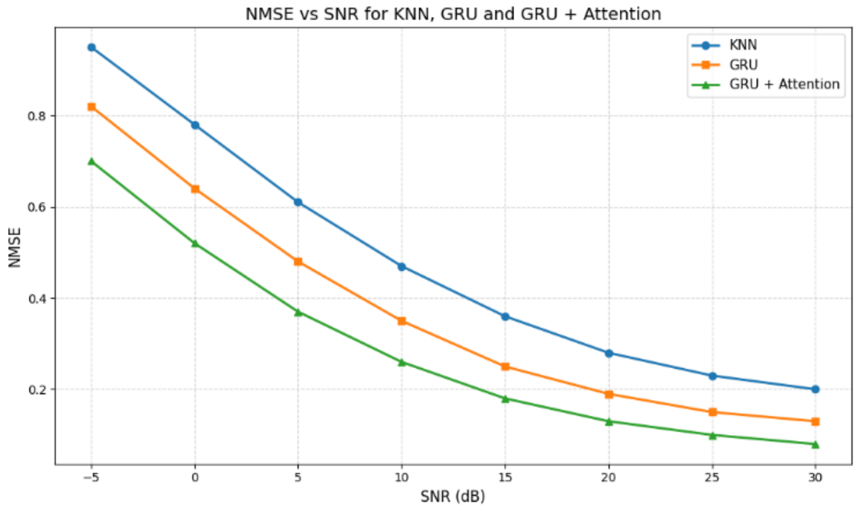

Fig. 1. NMSE vs SNR for KNN, GRU, and GRU with Attention

B. Confusion Matrix Analysis

Confusion matrices were also examined for each model, with channel observations split into Low, Medium, and High. Diagonal entries are correct classifications; everything off the diagonal is a misclassification between states. KNN's matrix (Fig. 2) shows moderate accuracy with a fair number of samples landing in the wrong bucket - not surprising, given how simple its decision rule is. GRU (Fig. 3) does better, since it can actually use relationships across the sequence, which cuts misclassifications in most states. GRU with attention (Fig. 4) improves further still, which points to the attention weighting helping the network pick out more useful features from the hybrid channel data.

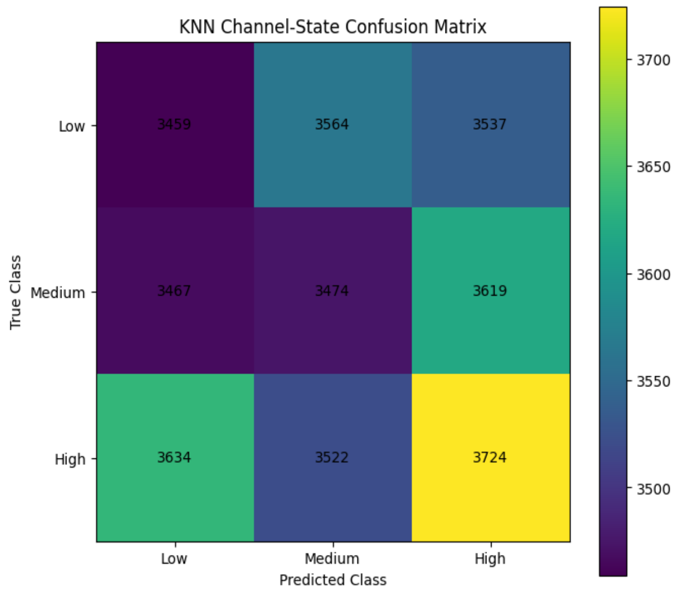

Fig. 2. KNN channel-state confusion matrix

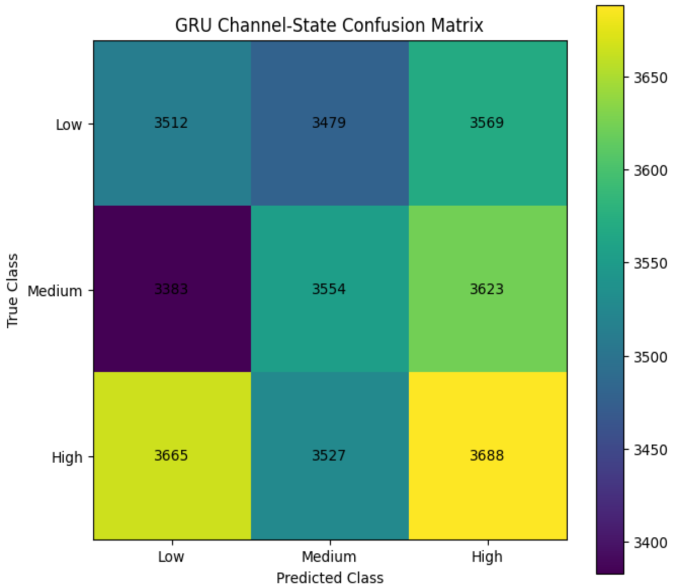

Fig. 3. GRU channel-state confusion matrix

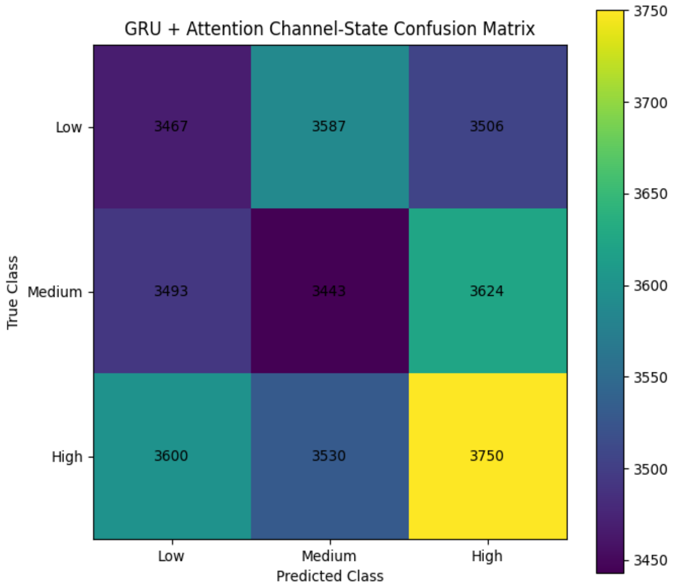

Fig. 4. GRU with Attention channel-state confusion matrix

C. Discussion

Two things stand out. First, accuracy improves with SNR for every model, which is exactly what would be expected as noise becomes less of a factor. Second, more complexity buys more accuracy here: KNN is the floor, GRU improves on it by learning sequential structure, and GRU with attention comes out on top by focusing on whatever features matter most. Taken together, this supports pairing recurrent learning with attention-based weighting for channel estimation in the kind of mixed near-field/far-field conditions RIS-assisted systems tend to produce.

V. CONCLUSION

An AI-based channel estimation framework has been built and tested here for RIS-assisted mmWave massive MIMO systems operating across both near-field and far-field propagation. The hybrid channel combined the direct base-station-to-user link with the RIS-reflected path, and OFDM-based pilots were used to generate the data needed for training and testing. Three estimators were compared along the way: KNN as a conventional baseline, GRU to pick up sequential dependencies in the channel data, and GRU with attention to refine that further by focusing on the most useful learned features.

Both NMSE and the confusion-matrix results point the same direction - the learning-based models beat the KNN baseline consistently, GRU improves on it through temporal-pattern learning, and GRU with attention comes out ahead of both by prioritising the most relevant features. This suggests deep-learning-based estimators are genuinely well suited to the added complexity that hybrid near-field/far-field propagation introduces, and that attention weighting adds something real over a plain recurrent model rather than just extra parameters.

Extending this to larger antenna and RIS configurations, more varied channel and mobility conditions, and the practical demands of real 5G/6G deployments is the natural next step.

REFERENCES

[1] T. C. Rapudu and O. O. Oyerinde, "Channel estimation for wideband multi-RIS-assisted mmWave massive-MIMO OFDM system with beam squint effect," IEEE Open J. Commun. Soc., vol. 6, pp. 4804-4817, Jun. 2025, doi: 10.1109/OJCOMS.2025.3575947.

[2] X. Wei and L. Dai, "Channel estimation for extremely large-scale massive MIMO: Far-field, near-field, or hybrid-field?," IEEE Commun. Lett., vol. 26, no. 1, pp. 177-181, Jan. 2022, doi: 10.1109/LCOMM.2021.3124927.

[3] J. Chen, Y.-C. Liang, H. V. Cheng, and W. Yu, "Channel estimation for reconfigurable intelligent surface aided multi-user mmWave MIMO systems," IEEE Trans. Wireless Commun., vol. 22, no. 10, pp. 6853-6869, Oct. 2023, doi: 10.1109/TWC.2023.3246264.

[4] S. Yang, W. Lyu, Z. Hu, and Z. Zhong, "Channel estimation for near-field XL-RIS-aided mmWave hybrid beamforming architectures," IEEE Trans. Veh. Technol., vol. 72, no. 8, pp. 11029-11034, Aug. 2023, doi: 10.1109/TVT.2023.3261340.

[5] M. Cui, Z. Wu, Y. Lu, X. Wei, and L. Dai, "Near-field communications for 6G: Fundamentals, challenges, potentials, and future directions," IEEE Commun. Mag., vol. 60, no. 6, pp. 54-60, Jun. 2022, doi: 10.1109/MCOM.004.2200136.

[6] M. A. ElMossallamy, H. Zhang, L. Song, K. G. Seddik, and Z. Han, "Reconfigurable intelligent surfaces for wireless communications: Principles, challenges, and opportunities," IEEE Trans. Cogn. Commun. Netw., vol. 6, no. 3, pp. 990-1002, Sep. 2020, doi: 10.1109/TCCN.2020.2992604.

[7] E. Björnson, H. Wymeersch, B. Matthiesen, P. Popovski, and L. Sanguinetti, "Reconfigurable intelligent surfaces: A signal processing perspective with wireless applications," IEEE Signal Process. Mag., vol. 39, no. 2, pp. 135-158, Mar. 2022, doi: 10.1109/MSP.2021.3130549.

[8] Y. Yang, B. Zheng, S. Zhang, and R. Zhang, "Intelligent reflecting surface meets OFDM: Protocol design and rate maximization," IEEE Trans. Commun., vol. 68, no. 7, pp. 4522-4535, Jul. 2020, doi: 10.1109/TCOMM.2020.2981458.

[9] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath, Jr., "Channel estimation and hybrid precoding for millimeter wave cellular systems," IEEE J. Sel. Topics Signal Process., vol. 8, no. 5, pp. 831-846, Oct. 2014, doi: 10.1109/JSTSP.2014.2334278.

[10] J. Lee, G.-T. Gil, and Y. H. Lee, "Channel estimation via orthogonal matching pursuit for hybrid MIMO systems in millimeter wave communications," IEEE Trans. Commun., vol. 64, no. 6, pp. 2370-2386, Jun. 2016, doi: 10.1109/TCOMM.2016.2557791.

[11] H. He, C.-K. Wen, S. Jin, and G. Y. Li, "Deep learning-based channel estimation for beamspace mmWave massive MIMO systems," IEEE Wireless Commun. Lett., vol. 7, no. 5, pp. 852-855, Oct. 2018, doi: 10.1109/LWC.2018.2832128.

[12] X. Wei, C. Hu, and L. Dai, "Deep learning for beamspace channel estimation in millimeter-wave massive MIMO systems," IEEE Trans. Commun., vol. 69, no. 1, pp. 182-196, Jan. 2021, doi: 10.1109/TCOMM.2020.3027027.

[13] H. Ye, G. Y. Li, and B.-H. F. 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.

[14] M. H. Rahman, M. A. S. Sejan, M. A. Aziz, and J.-I. Baik, "Deep learning-based improved cascaded channel estimation and signal detection for reconfigurable intelligent surfaces-assisted MU-MISO systems," IEEE Trans. Green Commun. Netw., vol. 7, no. 3, pp. 1515-1527, Sep. 2023, doi: 10.1109/TGCN.2023.3237132.

[15] A. Alkhateeb, "DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications," in Proc. Inf. Theory Appl. Workshop (ITA), San Diego, CA, USA, Feb. 2019.