Received 5 December 2025; revised 10 January 2026; accepted 17 January 2026. Date of publication 20 January 2026; date of current version 30 January 2026.

Digital Object Identifier 10.1109/OJCOMS.2026.3656443

AI-Driven Optimization of Wave-Controlled RISs via Sample-and-Hold Detection

GAL BEN-ITZHAK (Graduate Student Member, IEEE), ENDER AYANOGLU (Fellow, IEEE), FILIPPO CAPOLINO (Fellow, IEEE), AND A. LEE SWINDLEHURST (Life Fellow, IEEE)

Department of Electrical Engineering and Computer Science, University of California at Irvine, Irvine, CA 92697, USA

CORRESPONDING AUTHOR: E. AYANOGLU (e-mail: ayanoglu@uci.edu)

This work was supported in part by the National Science Foundation (NSF) under Grant 2030029.

ABSTRACT Reconfigurable Intelligent Surfaces (RISs) enhance wireless communication by steering electromagnetic waves through tunable reflection elements. However, scaling conventional RIS architectures to large arrays introduces significant wiring and control complexity. To address this, we investigate a wave-controlled RIS that uses biasing transmission lines carrying biasing standing waves (BSWs), which are sampled at each element to generate control voltages. We propose a sample-and-hold (S/H) circuit for precise temporal voltage acquisition and evaluate its performance in scenarios where neither channel state information (CSI) nor mathematical RIS models are available. Instead, optimization is fully data-driven: received power measurements are used to train a neural network (NN) that estimates the radiation pattern. A Genetic Algorithm (GA) refines the NN architecture, while Simulated Annealing (SA) tunes the BSW amplitudes to maximize a signal-to-leakage-plus-noise ratio (SLNR) metric using NN-based feedback. Optimized configurations are stored in a lookup table for real-time RIS control. We assess performance across diverse propagation conditions – including multipath and non-line-of-sight (NLoS) channels – under varying dataset sizes and SLNR targets. Relative to envelope-detector-based sampling, the proposed S/H circuit yields more accurate radiation pattern control and SLNR performance approaching that of an ideal system with perfect CSI and exact models. This idealized baseline further confirms that the S/H-based wave-controlled RIS can effectively optimize SLNR even in NLoS-dominated channels such as Rayleigh fading.

INDEX TERMS Machine learning (ML), simulated annealing (SA), genetic algorithm (GA), reconfigurable intelligent surface (RIS), sample-and-hold (S/H).

I. INTRODUCTION

RECONFIGURABLE Intelligent Surfaces (RISs) have emerged as a transformative technology for next-generation wireless networks, offering enhanced coverage, improved spectral efficiency, and resilience to environmental blockages by enabling programmable control over the wireless propagation environment [1]. A particularly promising variant is the wave-controlled RIS architecture, which leverages biasing standing waves (BSWs) on a transmission line (TL) to indirectly control the reflection coefficients of passive metasurface elements [2], [3], [4]. The combined BSWs are sampled at each RIS element location using dedicated circuitry, and converted into DC voltages used to bias varactor diodes and thus tune the reflection coefficients of the RIS elements. This allows for a reduced control overhead by modulating a smaller set of BSWs rather than controlling all RIS elements, as well as decreased wiring complexity by biasing entire rows of RIS elements with a single TL.

It was shown in [3] that the wave-controlled RIS can achieve strong beamforming and nullforming performance for two sampling circuits: an envelope detector and a sample-and-hold (S/H) circuit. To provide a practical evaluation of the wave-controlled RIS performance, the models used in [3] take into account the physical limitations of the varactor diodes, as well as biasing-dependent reflection magnitude losses, reductions in the dynamic range of the reflection phases, electromagnetic coupling between elements, and the

use of reduced sets of BSWs for RIS control. The algorithms described in that paper utilize a precise mathematical model of the RIS and assume perfect channel state information (CSI), as well as available live feedback from the receiver to optimize the BSWs and reflection coefficients under pure line-of-sight (LoS) channels. However, in realistic deployments, these assumptions break down: CSI acquisition is difficult or impossible for large passive RISs, live receiver feedback may not be available, and accurate EM-based modeling cannot be performed during operation [5], [6]. These constraints make online learning-based or other real-time adaptive optimization methods impractical for the wave-controlled RIS architecture considered here.

To address these problems, the follow-up work in [7] introduces an Artificial Intelligence (AI)-driven framework for optimizing wave-controlled RIS configurations without relying on precise modeling of the RIS or the environment and without requiring CSI or phase-shift design codebooks, specifically for the envelope detector-based architecture. Although the envelope detector is simple to implement due to its passive components, its nonlinear behavior and peak-detection constraints complicate RIS control. In this work, we extend the AI-driven framework to the S/H circuit, which provides improved linearity and stability, and we evaluate its performance under significantly more challenging propagation conditions by introducing Rician and Rayleigh NLoS channels. This allows us to assess the robustness of the wave-controlled RIS architecture beyond the pure LoS scenarios considered in prior work.

To ensure a fair and controlled comparison with the rectifier-based RIS architecture in [7], we deliberately retain the same optimization algorithms – a Genetic Algorithm (GA) for neural network (NN) architecture tuning and Simulated Annealing (SA) for RIS configuration optimization. The goal of this paper is not to introduce new optimization techniques, but to isolate the impact of the underlying circuit architecture on achievable performance across various channel models. The system constraints we consider – no CSI, no live receiver feedback, and no real-time EM modeling – also make online learning-based methods such as deep reinforcement learning (DRL) impractical, further motivating the use of offline, data-driven optimization [8], [9], [10]. In contrast, existing DRL-based RIS approaches typically assume either continuous interaction with the environment, real-time feedback, or access to CSI or reward signals during training – assumptions that are incompatible with the passive, feedback-limited wave-controlled RIS architecture studied here and with our requirement that optimization should be performed entirely offline.

In addition to the AI-driven approach, we also evaluate a baseline S/H-based RIS configuration under idealized assumptions where a closed-form RIS model, perfect CSI, and live receiver feedback are all available. This baseline illustrates the upper-bound performance achievable with the S/H circuit and provides a reference point against which the AI-driven framework can be compared across LoS, Rician, and Rayleigh channels.

The AI-driven framework is summarized as follows: We model the RIS-assisted wireless environment using an NN trained on received power measurements collected from the environment, enabling offline optimization of RIS configurations to achieve desired radiation patterns without requiring CSI acquisition or live receiver feedback. The process begins by exciting the RIS with random combinations of BSWs, sampling the resulting signal power at receivers positioned in target directions, and training the NN to learn the input-output relationship between the BSW amplitudes and corresponding radiation patterns. Optimization is performed using heuristic algorithms in two stages: First, the NN architecture is tuned via a GA to minimize validation loss and ensure a generalized and accurate mapping between inputs and outputs. Subsequently, SA is employed for power maximization and minimization tasks to optimize a signal-to-leakage-plus-noise ratio (SLNR) metric. It leverages the trained NN to infer the received powers at directions of interest as functions of different BSW amplitudes and iteratively converge towards an optimal setting, doing so offline instead of via live receiver feedback. During this process, optimal solutions are adaptively stored in a lookup table to enable rapid RIS configuration in future deployments or to serve as initialization points for generating more complex radiation patterns via SA. These algorithms are chosen not for novelty, but because they match the methodology of [7], enabling a direct and unbiased comparison between the S/H-based and rectifier-based RIS architectures.

We summarize the main contributions of the paper with the following four key points:

- *Circuit Model Modification:* In [7], we employed an envelope detector circuit using a half-wave rectifier to convert the composition of the BSWs on the TL at each RIS element location into a DC voltage. Here, we use the S/H circuit for voltage sampling. While more complex to implement, the S/H circuit offers improved linearity and stability, enabling more flexibility in control over the reflection coefficients [3].
- *Dataset Scaling and Diversity:* We investigate the impact of training dataset size on NN training and generalization performance. We consider the tradeoffs between increased model accuracy using larger datasets versus reduced training and dataset collection times with smaller datasets.
- *Channel Model Variation:* To assess the robustness of the proposed framework, we simulated the RIS-assisted wireless environment under diverse and more realistic propagation scenarios, from wireless channels characterized purely as LoS, to Rician fading with multipath scattering and NLoS components, to purely NLoS Rayleigh fading channels.
- *Baseline Case:* We evaluate a baseline case of the S/H-based wave-controlled RIS to optimize the SLNR

metric with the presence of a closed-form expression of the RIS model, perfect CSI, and available feedback from the receiver for iterative optimization. We compare the performance of this baseline against the AI-driven approach across multiple channel models and SLNR optimization targets.

These generalizations not only validate the flexibility of the wave-controlled RIS architecture, but also demonstrate the broader applicability of the AI-based optimization framework to other RIS implementations and wireless conditions. By decoupling the optimization process from specific circuit models and channel assumptions, our approach paves the way for scalable, hardware-aware RIS control strategies that can adapt to dynamic environments and evolving system requirements.

The remainder of the paper is organized as follows. Section II presents the physical model of the RIS including equivalent circuit representations, and the assumed mathematical models for the signals and wireless channel. Section III describes the optimization algorithms in detail. Section IV provides simulation results for various scenarios and analyzes the performance of our RIS design approach. Finally, Section V compares the capabilities of this approach with an ideal scenario where the system is perfectly modeled.

II. MODELS AND ASSUMPTIONS

A. RIS MODEL

Our work in [2], [3], [7] and the physical experimental realization in [4] assume an RIS composed of a uniform linear array of M metallic patches separated by d_x along the x direction, each connected to a reverse-biased varactor diode. A biasing TL stretches along the array on the backside of the RIS and is terminated via a short circuit. The periodic time-domain signal generated by the waveform generator and carried by the TL is composed of a set of N BSWs at harmonic frequencies nf_b, with fundamental frequency $f_b = v_{ph}/\lambda_b$, where λ_b is the guided wavelength projected along the x direction, and $v_{ph} = c/n_{slow}$ is the phase velocity of the propagated modes projected along x. The slowness factor $n_{slow} = (L_p/d_x)\sqrt{\epsilon_{eff}}$ is determined by the geometry of the biasing TL, where L_p is the curvilinear distance along the TL between adjacent RIS elements (Fig. 2), and ϵ_{eff} is the effective permittivity that accounts for material properties [3], [4], [7]. The composite BSW signal along the TL is given by

$$w(x,t) = W_0 + \sum_{n=1}^{N} W_n \sin\left(\frac{n\pi\,(x + L_{left})}{L_{tot}}\right)\sin(n\omega_b t) \quad (1)$$

where W_0 is the DC offset for varactor biasing, W_n is the tunable amplitude of the n-th harmonic, L_{left}, L, and L_{right} are TL segment lengths along the x-axis before, between, and after the first and last unit cells, respectively, such that $L_{tot} = L_{left} + L + L_{right}$. In this paper we have chosen that $L_{tot} = \lambda_b/2$, so that the $n = 1$ standing wave is half period of a sine as in Fig. 1, though other lengths could be chosen

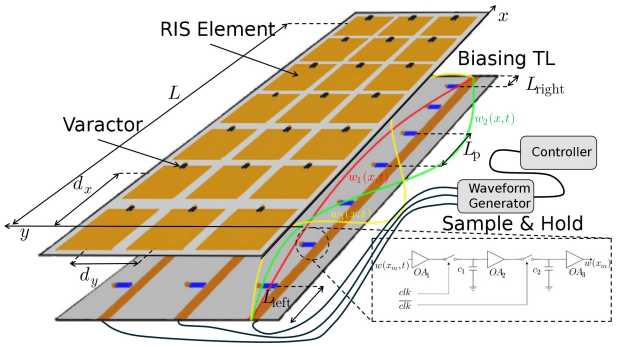

FIGURE 1. Full view of the wave-controlled RIS. The top layer comprises RIS elements as metallic patches on a grounded substrate, connected to grounded varactors, uniformly separated by d_x and d_y along the x and y directions, respectively. Each patch is biased by a DC voltage through a vertical via to the bottom layer (the backside of the surface), which contains biasing transmission lines (TLs), each carrying N BSWs. DC biasing is achieved by sampling the BSW composition at each unit cell location using S/H circuits. The BSWs amplitudes and frequencies are determined by a periodic time domain signal injected from a waveform generator connected to a programmable external controller. On the bottom layer, the adjacent sampling circuits are separated by d_x. The length of the TL between the connector and the first unit cell is L_{left}, the segment between the last unit cell and the TL termination is L_{right}, and the segment in-between is L.

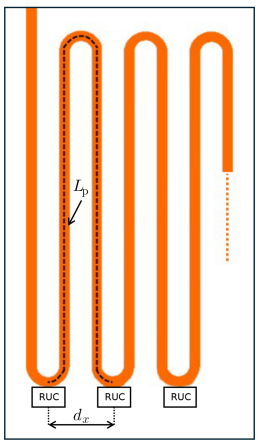

FIGURE 2. Serpentine geometry of the biasing TL made by a microstrip on a grounded substrate (not shown here), placed under the RIS metallic patches on the top layer of Fig. 1. The linear distance (along x) between adjacent RIS Unit Cells (RUCs), each one containing the S/H circuit and a biasing via to connect to a varactor on the top layer, is denoted by d_x, while the curvilinear distance between them along the TL is L_p.

as well. The biasing TL under the RIS patches is made of the serpentine microstrip waveguide in Fig. 2 over a grounded substrate, Rogers RT5880LZ, with relative permittivity $\epsilon_r = 2$ and dielectric loss $\tan(\delta) = 0.0021$. The varactors used to control the RIS elements directly on the top layer of Fig. 1 are SMV1231-040LF, provided by Skyworks Solutions, Inc.

Each RIS element samples the waveform at discrete positions $x_m = md_x$ using the S/H circuit [3]. This circuit is commonly used in communication systems and analog-to-digital converters (ADCs) to sample a time-varying signal and hold it as a constant voltage for digitization [11]. A general implementation of the S/H circuit with a main-secondary architecture and a sampling clock signal is shown in the lower right portion of Fig. 1 and the operation of the circuit is demonstrated in Fig. 3. The S/H initially captures the input signal using a buffer with high input impedance,

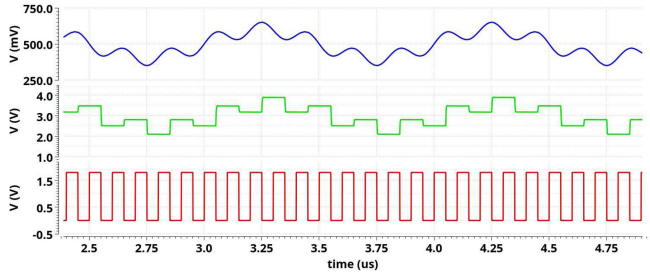

FIGURE 3. Operation of the proposed sample-and-hold circuit, sampling a superposition of sinusoids over a few microseconds. Top: Input waveform. Center: Sampled waveform. Bottom: Control signal.

such that its effects on the signal are minimal. When the circuit is in the "sample" mode, the control signal (*clk*) is high and capacitor C_1 tracks the input waveform as the switch after OA_1 closes to charge C_2 to the same voltage as C_1. When the circuit is in the "hold" mode, the control signal is low, the switch after OA_1 opens such that the voltage on C_1 is held constant, and the switch between OA_2 and C_2 closes. C_2 then holds a constant voltage which then gets buffered through OA_3 to the output. Essentially, the output DC signal would change on the falling edge of the control signal.

Compared to the envelope detector, the S/H circuit provides superior control over the RIS reflection behavior, as well as improved linearity by minimizing mismatches and losses across the TL, without relying on a time constant for peak detection. Additional DC gain at the buffer stage also enables lower BSW amplitudes, relaxing the requirements on the waveform generator. These benefits come at the cost of increased hardware complexity: each RIS element requires active sampling switches, small hold capacitors, and buffers, along with routing for biasing and sampling signals on metal layers beneath the surface. This overhead corresponds to an area and power cost on the order of tens of transistors per element and scales linearly with the number of RIS elements. The sampling signal distribution network likewise scales linearly with the number of rows, but operates at low frequencies (kHz-MHz), which relaxes jitter and skew constraints. Because the S/H circuit samples slowly varying BSW envelopes rather than RF waveforms, timing mismatch does not accumulate with array size. With careful design to isolate sampling signal interference, potentially a single wire traversing all elements can be employed. The additional hardware complexity remains manageable for large RISs and helps contextualize the significant SLNR improvements demonstrated in this work.

We assume that all S/H circuits sample the BSW composition at the same time t_0, after which it is held constant as a DC voltage used as reverse bias on each varactor. The DC bias at each varactor indexed by m is given by

$$w(m) = w(md_x, t_0) = W_0 + \sum_{n=1}^{N} W_n \sin\left(\frac{n\pi\,(md_x + L_{\text{left}})}{L_{\text{tot}}}\right) \sin(n\omega_{\text{b}} t_0). \quad (2)$$

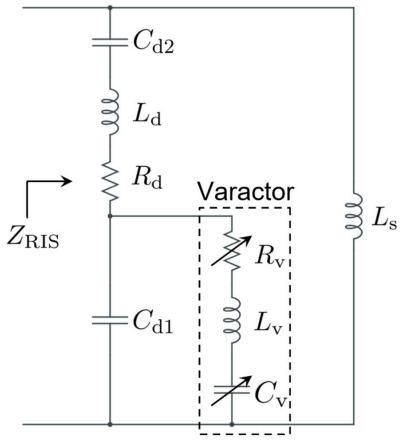

FIGURE 4. Analytical model to evaluate the RIS reflection coefficient Γ, accounting for the presence of the varactor, the geometry of the patch and substrate, losses and coupling effects.[1] The varactor components are voltage-dependent. The equivalent impedance seen by the incoming plane wave is Z_{RIS}.

Reflection from RIS unit cells is modeled using a passive equivalent circuit that includes the varactor model and parasitics from the pads, package, substrate, as well as interactions (i.e., electromagnetic coupling) with other metallic patches in the periodic structure, as shown in Fig. 4 [4], [12], [13].

Using full-wave simulations based on the finite element method (FEM) implemented in CST Microwave Studio to model an array of RIS elements (accounting for their couplings), we find the values for the equivalent circuit elements used to fit the simulated data: metallic patch parameters $R_{\text{d}} = 0.17\ \Omega$ and $L_{\text{d}} = 1.92$ nH; fringe and substrate capacitances associated with the metallic patch on the varactor side $C_{\text{d1}} = 1.25$ pF; fringe and substrate capacitances due to other side of the patch $C_{\text{d2}} = 2.72$ pF; inductance from the grounded substrate due to the short circuit termination of the transmission line $L_{\text{s}} = 1.60$ nH; and inductance from varactor package, pads, and connections $L_{\text{v}} = 2.39$ nH. The varactor impedance is

$$Z_{\text{v}}(V) = R_{\text{v}}(V) + j\omega L_{\text{v}} + \frac{1}{j\omega C_{\text{v}}(V)} \quad (3)$$

and the total impedance is

$$Z_{\text{RIS}} = \left(R_{\text{d}} + j\omega L_{\text{d}} + \frac{1}{j\omega C_{\text{d2}}} + Z_{\text{v}} \| \frac{1}{j\omega C_{\text{d1}}}\right) \| j\omega L_{\text{s}}. \quad (4)$$

The reflection coefficient is computed as

$$\Gamma = \frac{Z_{\text{RIS}} - Z_0}{Z_{\text{RIS}} + Z_0}, \quad (5)$$

where Z_0 is the free-space impedance. In Fig. 4, $R_{\text{v}}(V)$, $C_{\text{v}}(V)$ are the varactor parameters obtained from the vendor's information, which depend on the reverse biasing voltage V across its terminals, plotted in Fig. 5(a) [14]. The

[1]This figure replaces [7, Fig. 4], where the splitting of the capacitance C_{d} into two separate capacitors, C_{d1} and C_{d2}, was omitted from the figure but accounted for in the numerical results.

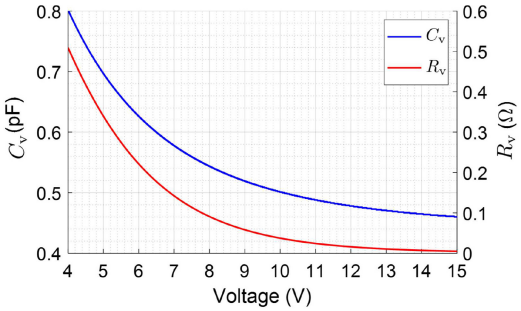

(a) Equivalent resistance and capacitance characteristics of the Skyworks SMV1231-040LF varactor [7].

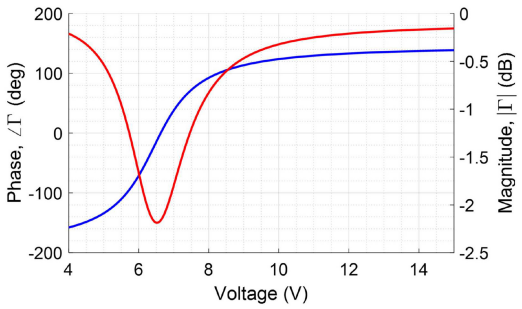

(b) Reflection characteristics (magnitude and phase) of the RIS unit cell [7].

FIGURE 5. RIS unit cell characteristics at $f_c = 2.45$ GHz, showcasing the tunability of the passive varactor diode as a function of reverse biasing voltage and the corresponding RIS reflection coefficients.

magnitude and phase of the resulting reflection coefficient for a carrier frequency $f_c = 2.45$ GHz are plotted over the allowable range of reverse voltage bias [4 V, 15 V] in Fig. 5(b).

B. SIGNAL AND CHANNEL MODELS

The propagation environment consists of a single-antenna basestation (BS) and K single-antenna users. The wireless far-field channels are characterized by a flat fading Rician channel model

$$\boldsymbol{h}(\theta) = \sqrt{\frac{\kappa}{\kappa+1}}\boldsymbol{h}_{\text{LoS}}(\theta) + \sqrt{\frac{1}{\kappa+1}}\boldsymbol{h}_{\text{NLoS}} \tag{6}$$

with Rician factor κ and azimuth angle θ [15]. The LoS channels are each expressed as the array response $\boldsymbol{h}_{\text{LoS}} = [1, e^{j\omega(\theta)}, \ldots, e^{j(M-1)\omega(\theta)}]^T$ with spatial frequencies $\omega(\theta) = 2\pi d_x(f_c/c)\sin(\theta)$. When $\kappa = 0$, only NLoS paths exist and the channels are characterized by Rayleigh fading. The elements of the NLoS channels are distributed as standard circularly symmetric complex Gaussian random variables $\boldsymbol{h}_{\text{NLoS}} \in \mathcal{CN}(0,1)^{M\times 1}$. We assume that the RIS orientation and the UE angular region of interest remain approximately fixed between dataset collection and radiation pattern optimization. This ensures that the dominant LoS steering directions are consistent, while allowing the NLoS components to vary.

The received baseband signal at the k-th user equipment (UE) is given by

$$y_k = \sqrt{P}\boldsymbol{h}_{\text{BS}}^T\boldsymbol{\Gamma}\boldsymbol{h}_k s_k + n_k, \tag{7}$$

where s_k is the transmitted signal, P is the transmit power, $n_k \in \mathcal{CN}(0,\sigma_n^2)$ is Additive White Gaussian Noise (AWGN), and $\boldsymbol{h}_{\text{BS}}$ and $\boldsymbol{h}_k$ are the $M \times 1$ vectors representing the channel between the BS and the RIS, and between the RIS and the k-th UE, respectively [6], [16]. For a UE at azimuth angle θ_k from the RIS, $\boldsymbol{h}_k = \boldsymbol{h}(\theta_k)$ as in (6). To assess the performance gained by employing the RIS, we assume that there is no direct link between the BS and each UE. The BS is located at the broadside of the RIS, creating normal incidence such that $\boldsymbol{h}_{\text{BS}} = \boldsymbol{h}(0)$. The diagonal matrix $\boldsymbol{\Gamma}$ contains the RIS reflection coefficients, given by

$$\boldsymbol{\Gamma} = \text{diag}[\Gamma(0), \Gamma(1), \ldots, \Gamma(M-1)], \tag{8}$$

where $\Gamma(m)$ is the reflection coefficient at the m-th RIS element, and due to passivity constraints, $|\Gamma(m)| \leq 1$ for all $m = 0, 1, \ldots, M-1$.

The signal-to-noise ratio (SNR) at each UE k is defined as the ratio of the received signal power to the noise power σ_n^2,

$$\text{SNR}_k = \frac{|E[y_k]|^2}{\sigma_n^2} = \frac{|E\left[\sqrt{P}\boldsymbol{h}_{\text{BS}}^T\boldsymbol{\Gamma}\boldsymbol{h}_k s_k\right]|^2}{\sigma_n^2} = \frac{P|\boldsymbol{h}_{\text{BS}}^T\boldsymbol{\Gamma}\boldsymbol{h}_k|^2}{\sigma_n^2}. \tag{9}$$

We define an additional metric, signal-to-leakage-plus-noise ratio (SLNR) for $K \geq 1$ desired UE and $L \geq 0$ undesired locations as

$$\text{SLNR} = \frac{\min_{i\in\{1,2,\ldots,K\}} P|\boldsymbol{h}_{\text{BS}}^T\boldsymbol{\Gamma}\boldsymbol{h}_{\text{d},i}|^2}{\max_{j\in\{1,2,\ldots,L\}} P|\boldsymbol{h}_{\text{BS}}^T\boldsymbol{\Gamma}\boldsymbol{h}_{\text{e},j}|^2 + \sigma_n^2}, \tag{10}$$

where each $\boldsymbol{h}_{\text{d},i}$ represents the channel from the RIS to the i-th desired UE and each $\boldsymbol{h}_{\text{e},j}$ represents the channel from the RIS to the j-th undesired location. The null directions may correspond to UE directions that require low signal interference for better quality of service. Note that for $K = 1$ and $L = 0$, the SLNR expression reduces to the SNR expression in (9). The SLNR metric is used to analyze the performance of the RIS by evaluating the resulting radiation patterns and can be generalized to maximize or minimize received powers in given directions of interest. In this work, we seek to maximize the SLNR metric for different RIS configurations and channels using the S/H-based wave-controlled RIS. This is achieved by searching for an optimal set of BSW amplitudes $\boldsymbol{W} = [W_1, W_2, \ldots, W_N]^T$ using the algorithms discussed in the following sections. We emphasize that the channel models discussed in this section are used to simulate the RIS-aided wireless environment for evaluation of the SLNR metric, but are not used as part of optimization in Sections III and IV.

III. ALGORITHMS

A. DATASET GENERATION

In practice, the duration of dataset collection depends on the desired dataset size $\mathcal{D}$, the number of BSW amplitudes N, and the number of UE directions n_{angles} at which measurements are acquired. For each BSW configuration, the receivers at the selected directions record the received power once the system reaches steady state, so the total collection time scales linearly with $\mathcal{D} \cdot n_{\text{angles}}$. As our results in Table 3 show, the required dataset size can be reduced significantly when only coarse beam steering is needed or when the angular region of interest is limited to a small set of UE directions. Consequently, dataset collection can be performed offline over a timescale that is compatible with slowly varying deployments where the RIS orientation and UE regions of interest remain approximately fixed.

In our approach, we design an NN that models the entire RIS-aided wireless environment, mapping BSW amplitudes to received power values in a radiation pattern. The procedure to generate the training dataset for the NN using the S/H-based RIS is outlined in Algorithm 1 and summarized below. Each element in the dataset represents a pair of input and output arrays, where the input array $\boldsymbol{W}$ contains the N BSW amplitudes and the output array $\boldsymbol{P}$ contains samples of the received signal power at n_{angles} discrete UE directions. In simulation, this power value is calculated by the numerator of (9), though in the practical case, the receivers at each direction of interest would measure and provide the received signal power.

We begin by initializing each element of the N-dimensional array with BSW amplitudes drawn from a Gaussian distribution with small variance σ_1^2. Then, we randomly select a subset $\boldsymbol{S}$ of BSW amplitudes from each N-dimensional array and replace the elements in $\boldsymbol{S}$ with values drawn from a Gaussian distribution with larger variance σ_2^2. This method excites a small set of BSWs with higher voltages and treats the remainder as "noise" to assist the NN during training and avoid overfitting. This becomes important later when SA with feedback from the trained NN is used to optimize the radiation patterns, since SA operates by making small random perturbations to each input array element, passing the inputs through the NN, and repeatedly calculating the objective function from the inferred NN output until it converges to a local optimum. Thus, the NN learns to distinguish between the BSW frequencies that significantly contribute to the shape of the radiation pattern and those that do not. We verify that the resulting DC voltages from the sampled BSWs do not exceed the varactor voltage bias limits. Then, we sample the resulting received power (in dB) at all n_{angles} UE directions by simulating the wireless channels and the excited RIS and store the arrays in memory. The synthesis of the $\boldsymbol{W}$ arrays and their corresponding $\boldsymbol{P}$ arrays is repeated $\mathcal{D}$ times until the dataset is full.

During dataset generation in our simulation, the wireless channel is held constant to ensure that the NN learns a stable mapping between the BSW amplitudes and the resulting radiation pattern. This does not restrict the applicability of the model to dynamic environments. In practice, the dominant LoS component of the RIS-UE link remains fixed as long as the RIS orientation and the UE directions of interest do not change. The NN therefore learns the underlying geometric relationship between the BSW excitations and the far-field pattern, rather than memorizing a specific channel realization. As shown later in our results, the trained NN generalizes to scenarios with varying NLoS components – including Rician and Rayleigh fading – by effectively filtering out the random multipath contributions and forming beams based on the stable LoS structure.

Algorithm 1 Dataset Generation

1: $i \leftarrow 0$
2: **while** $i < \mathcal{D}$ **do**
3: $\quad \boldsymbol{W}(n) \leftarrow$ sample from $\mathcal{N}(0, \sigma_1^2)$ for $n = 1, \ldots, N$.
4: $\quad k \leftarrow$ integer sample from $\mathcal{U}(1, N)$.
5: $\quad \boldsymbol{S} \leftarrow$ set of k uniquely and uniformly selected integers from 1 to N.
6: $\quad \boldsymbol{W}(n) \leftarrow$ sample from $\mathcal{N}(0, \sigma_2^2)$ $\forall n \in \boldsymbol{S}$.
7: $\quad$ Calculate maximum absolute value $v_{\max}$ of the resulting BSW superposition when sampled by the S/H at t_0 from W_0 and $\boldsymbol{W}$ using (2).
8: $\quad$ **if** $v_{\max}$ is outside varactor biasing bounds **then**
9: $\quad\quad$ Go back to line 3.
10: $\quad$ **end if**
11: $\quad$ Sample the power (in dB) at all n_{angles} desired directions, store in $\boldsymbol{P}$.
12: $\quad$ Store $\boldsymbol{W}$ and $\boldsymbol{P}$ as $\boldsymbol{W}_i$ and $\boldsymbol{P}_i$.
13: $\quad i \leftarrow i + 1$
14: **end while**
15: Find the maximum absolute value s_{w} from all $\boldsymbol{W}_i$ in the dataset, $i = 0, 1, \ldots, \mathcal{D}$.
16: Find the maximum absolute value s_{p} from all $\boldsymbol{P}_i$ in the dataset, $i = 0, 1, \ldots, \mathcal{D}$.
17: Normalize $\boldsymbol{W}_i \leftarrow \boldsymbol{W}_i / s_{\text{w}}, \quad i = 0, 1, \ldots, \mathcal{D}$.
18: Normalize $\boldsymbol{P}_i \leftarrow \boldsymbol{P}_i / s_{\text{p}}, \quad i = 0, 1, \ldots, \mathcal{D}$.

After the datasets are generated and before training the NN model, all $\boldsymbol{W}$ and $\boldsymbol{P}$ arrays in each dataset are normalized to the range $[-1, 1]$ by dividing them by the maximum absolute BSW amplitude in the dataset, s_{w}, and the maximum absolute dB power in the dataset, s_{p}.

B. NEURAL NETWORK ARCHITECTURE OPTIMIZATION USING A GENETIC ALGORITHM

We model the wireless system using the NN depicted in Fig. 6 as a nonlinear mapping $\hat{f}(\boldsymbol{W}) = \hat{\boldsymbol{P}}$, where $\boldsymbol{W} \in \mathbb{R}^N$, and $\hat{\boldsymbol{P}} \in \mathbb{R}^{n_{\text{angles}}}$ approximates the true radiation pattern $\boldsymbol{P}$. The NN hyperparameter optimization is difficult due to the high-dimensional, non-convex search space and the complex interdependencies among parameters such as hidden layer count, number of nodes per layer, activation functions, and training settings [18]. Traditional methods such as grid search

TABLE 1. Summary of parameters for the GA and SA.

Algorithm Name	Parameter Description	Symbol	Value
Genetic Algorithm	Number of Epochs	-	$\{125, 126, \ldots, 250\}$
	Batch Size	-	$\{32, 64, 128, 256, 512\}$
	Number of Hidden Layers	-	$\{1, 2, \ldots, 7\}$
	Number of Nodes per Layer	-	$\{64, 128, 256, 512, 1024, 2048\}$
	Activation Functions [17]	-	{ReLU, PReLU, Sigmoid, Tanh}
	Initial Population Size	N_{pop}	32
Simulated Annealing	Cooling Factor	k_{c}	0.002
	Number of Iterations	$i_{\max}$	2000
	Step Size	λ	0.015
	Maximum Number of Iterations Between Current SLNR and Previous Best SLNR	i_{rst}	200
	Temperature Scaling Factor	T_{scale}	100

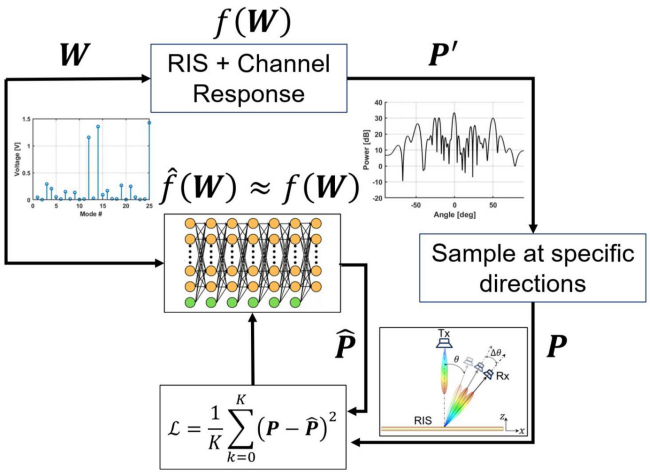

FIGURE 6. The system is modeled by a closed box function $f(W)$ which encapsulates the combined effects of the RIS and the wireless channels. This function takes as input a set of BSW amplitudes W and produces a corresponding radiation pattern P'. The radiation pattern is then sampled at predefined spatial locations to yield the observed power values P. An NN is trained to estimate the sampled power values $\hat{P}$ from the same input data and estimate the mapping function as $\hat{f}(W) \approx f(W)$. The training objective minimizes the MSE between the true and estimated power values, defined as the loss function $\mathcal{L}$ [7].

or gradient descent often require extensive computation or fail to escape local optima.

Instead, we employ a Genetic Algorithm (GA), a global optimization method well-suited for estimating discrete and mixed hyperparameters [19]. GA maintains a diverse population and uses crossover and mutation to efficiently explore the search space without relying on gradient information, evaluating each individual by its fitness as defined by a desired cost function. Originally applied to optimize NNs for classification tasks in [20], GA is also effective for regression tasks when an appropriate cost function is defined (in this case, validation loss), and the genetic operators are adapted to accommodate varying hyperparameter sets. This approach automates NN design, improving performance and eliminating manual tuning. In [7], we demonstrated the utility of GA in optimizing the NN architecture for an envelope detector-based, wave-controlled RIS with LoS channels and $N = 25$ BSWs. In this work, we extend the GA framework to S/H-based RIS under diverse channel models, varying training dataset sizes, and different BSW counts, illustrating its applicability for broader RIS optimization scenarios.

Pseudocode for the GA is given in Algorithm 2. The algorithm begins by initializing a population of NNs with randomly chosen hyperparameters from Table 1. The NNs are trained on the dataset containing the BSWs and power measurements, and their validation losses are recorded. From the population, a subset of "parent" NNs are randomly chosen to create two new "children" NNs, based on principles of crossover and mutation. The children NNs are also trained on the same dataset. The population size is then reduced, keeping only the NNs with the smallest validation losses. This evolution continues until the population size is one and the NN with the smallest validation loss remains.

C. BEAM PATTERN OPTIMIZATION USING SIMULATED ANNEALING AND LOOKUP TABLE

Using the trained NN, offline optimization of the BSW amplitudes $\boldsymbol{W}$ is performed to configure the RIS for a target radiation pattern. Simulated Annealing (SA) is employed for this task, leveraging prior success in SLNR-based optimization for wave-controlled RIS using envelope detectors [3], [7]. The NN estimates directional power samples after each update of $\boldsymbol{W}$, guiding SA toward improved SLNR.

Algorithm 2 Genetic Algorithm for NN Optimization

Initialize a population of N_{pop} NNs each with hyperparameters randomly chosen from Table 1.
Train the NNs on the training dataset, record their validation losses.
while $N_{\text{pop}} > 1$ **do**
    Randomly select $\frac{N_{\text{pop}}}{2}$ pairs of parents.
    Each pair of parents creates 2 children with the following principles:
    Crossover: Number of epochs, batch size, and number of hidden layers of each child are chosen randomly between the parents.
    if child has the same number of hidden layers as the parent with the smaller number of hidden layers **then**
        Activation functions and number of nodes are each randomly inherited from either parent for all hidden layers.
    else
        Activation functions and numbers of nodes are each randomly inherited from either parent for the smaller number of hidden layers between the parents. The parameters of the additional layers are selected randomly either from the last hidden layer of the parent with the smaller number of layers, or the corresponding layers from the parent with the greater number of layers.
        Mutation step: Randomly shuffle the numbers of nodes and activation functions between all hidden layers of the child.
    end if
    Train all child NNs, record their validation losses.
    Keep the $\frac{N_{\text{pop}}}{2}$ individuals from the entire population with the smallest validation losses.
    $N_{\text{pop}} \leftarrow \frac{N_{\text{pop}}}{2}$
end while
return NN with smallest validation loss.

Algorithm 3 Simulated Annealing With Neural Network

Initialize $\boldsymbol{W} \leftarrow \boldsymbol{0}$, $\boldsymbol{W}_{\text{best}} \leftarrow \boldsymbol{W}$, $i_{\text{best}} \leftarrow 0$.
Foward pass $\boldsymbol{W}$ to the NN, scale the output using (13), calculate $\text{SLNR}_{\text{current}}$ with (12), and set $\text{SLNR}_{\text{best}} \leftarrow \text{SLNR}_{\text{current}}$.
for $i = 1$ to $i_{\max}$ **do**
    if $i - i_{\text{best}} \geq i_{\text{rst}}$ **then**
        Reset $\boldsymbol{W} \leftarrow \boldsymbol{W}_{\text{best}}$, $i_{\text{best}} \leftarrow i$ $\text{SLNR}_{\text{current}} \leftarrow \text{SLNR}_{\text{best}}$.
    end if
    $T \leftarrow T_{\text{scale}}(1 - i/i_{\max})$.
    Perturb: $\boldsymbol{W}_{\text{new}}(n) \leftarrow \boldsymbol{W}(n) + \mathcal{N}(0, \lambda^2)$ for $n = 1, \ldots, N$.
    Evaluate DC bias using (14).
    if bias exceeds limits **then**
        Continue to next iteration.
    end if
    Compute SLNR_{new} by NN forward pass of $\boldsymbol{W}_{\text{new}}(n)$, evaluate using (13) and (12).
    if $\text{SLNR}_{\text{new}} > \text{SLNR}_{\text{best}}$ **then**
        Update: $\boldsymbol{W}, \boldsymbol{W}_{\text{best}} \leftarrow \boldsymbol{W}_{\text{new}}$.
        $\text{SLNR}_{\text{current}}, \text{SLNR}_{\text{best}} \leftarrow \text{SLNR}_{\text{new}}$, $i_{\text{best}} \leftarrow i$.
    else
        Compute p from (11).
        if $p \geq \text{rand}(1)$ **then**
            Update: $\boldsymbol{W} \leftarrow \boldsymbol{W}_{\text{new}}$, $\text{SLNR}_{\text{current}} \leftarrow \text{SLNR}_{\text{new}}$.
        end if
    end if
end for

We adopt the SA algorithm from [7] to stochastically optimize $\boldsymbol{W}$ for SLNR maximization, outlined in Algorithm 3. Starting at high temperature T, the algorithm explores diverse configurations via random perturbations to $\boldsymbol{W}$, gradually cooling to favor SLNR-improving updates. The transition probability from $\boldsymbol{W}$ to $\boldsymbol{W}_{\text{new}}$ is defined as

$$p = \begin{cases} 1 & \text{SLNR}_{\text{new}} > \text{SLNR}_{\text{current}} \\ e^{\left(-\frac{\text{SLNR}-\text{SLNR}_{\text{new}}}{k_c T}\right)} & \text{SLNR}_{\text{new}} \leq \text{SLNR}_{\text{current}}, \end{cases} \tag{11}$$

where $k_c > 0$ is a cooling constant. This enables SA to escape local minima by exploring solutions that do not always improve the SLNR [21]. The SLNR is calculated using a similar expression to (10),

$$\text{SLNR} = \frac{\min_{i\in\{1,2,\ldots,K\}} P_{\theta_i}}{\max_{j\in\{1,2,\ldots,L\}} P_{\theta_j} + \sigma_n^2}, \tag{12}$$

where P_{θ_i} and P_{θ_j} are the target beam and null directions, respectively. Since the NN outputs dB power values scaled between $[-1, 1]$, we scale them back into power scale by modifying the NN output for a given direction θ_x as

$$P_{\theta_x} = 10^{s_p p_{\theta_x}/10}, \tag{13}$$

where p_{θ_x} is the NN output and s_p is the scaling factor corresponding to the maximum absolute dB power value in the dataset, as in Section III-A.

The SLNR optimization is further constrained by varactor biasing limits (4-15 V). Therefore, each $\boldsymbol{W}$ update is validated against these limits by calculating the corresponding DC voltages as in (2), serving as a conservative bound for biasing feasibility since the voltages may experience some nonlinear and frequency-dependent losses from the TL and sampling circuits. Note that the scaling of the $\boldsymbol{W}$ arrays when the dataset was created enforced them to the limits $[-1, 1]$ during NN training. Therefore, the $\boldsymbol{W}$ array is scaled by s_w to restore the original amplitude scaling, resulting in the modified expression

$$\begin{aligned} w(m) &= w(md_x, t_0) \\ &= W_0 + s_w \sum_{n=1}^{N} W_n \sin\left(\frac{n\pi\,(md_x + L_{\text{left}})}{L_{\text{tot}}}\right) \sin(n\omega_b t_0). \end{aligned} \tag{14}$$

SA is computationally efficient: perturbations involve simple arithmetic, and SLNR is evaluated via NN inference. The algorithm is initialized with $\boldsymbol{W} = \mathbf{0}$, unless a more informed starting point is available.

A lookup table of optimized BSW amplitude vectors $\boldsymbol{W}$ is leveraged to accelerate convergence and reduce the number of SA iterations. These stored solutions enable rapid RIS reconfiguration by reusing previously computed solutions. When optimizing more complex radiation patterns, the existing solutions for $\boldsymbol{W}$ may be used as initializations to improve SA efficiency for multi-beam or multi-null scenarios. For instance, a $\boldsymbol{W}$ that generates a beam at θ_1 can be used to initialize the design of a dual-beam pattern at both θ_1 and θ_2. This can be done by using the $\boldsymbol{W}$ solution for θ_1 instead of an array of empty BSW amplitudes in the first SA step, adjusting the objective function to account for both beam directions, and iterating from there to converge towards a new optimum. Similar approaches are used to generate multiple beams, as well as to form beams and then form nulls, by generating one beam at a time and adjusting the objective function to either account for more beams, or to account for nulls after all the beams have been created. The nullforming is left for the last optimization step due to the sensitivity of nulls to small changes in the reflection coefficients that significantly impact the SLNR-based objective function.

The lookup table does not require periodic refreshing as long as the RIS orientation and the set of UE directions of interest remain unchanged. Because the NN learns the geometric mapping between the BSW amplitudes and the far-field radiation pattern – rather than a specific instantaneous channel realization – the stored $\boldsymbol{W}$ solutions remain valid even as the NLoS components of the environment drift over time. As demonstrated in our results, the trained NN effectively filters out random multipath variations and forms beams based on the stable LoS structure. New entries are added to the table only when additional UE positions or SLNR targets are required. If the RIS orientation changes or the dominant LoS geometry shifts significantly, the mapping between BSW amplitudes and beam directions changes, and a new dataset must be collected to train the NN appropriately and regenerate the lookup table.

Given the large training dataset of random $\boldsymbol{W}$ vectors and their associated power samples, we construct the lookup table as follows. For a desired beam direction θ, we identify the $\boldsymbol{W}$ yielding maximum power $P_{\max}$ in that direction, either from the training dataset or from existing entries in the lookup table, and use it to initialize SA. The converged $\boldsymbol{W}$ is then stored along with its beam direction. For multiple beams, we select the strongest existing entry based on SLNR for any subset of the beams and refine it via SA to add additional beams, one at a time. All nulls are appended in the final optimization stage. Over time, the table adaptively expands, reducing future optimization overhead.

Each entry in the table includes:

- Beam directions: $\{\theta_{\text{b1}}, \theta_{\text{b2}}, \ldots\}$
- Null directions: $\{\theta_{\text{n1}}, \theta_{\text{n2}}, \ldots\}$
- SLNR [dB]
- BSW amplitudes $\boldsymbol{W}$

Additionally, we apply linear interpolation to estimate power at arbitrary angles between the uniformly sampled directions, where the samples are separated at most by the 3 dB beamwidth corresponding to the size of the RIS. This ensures that nearby directions within a beamwidth remain sensitive to power adjustments, which is critical for effective beam- and nullforming. Linear interpolation is used to find the value of the beam pattern for any unsampled target angles. This approach allows smooth SLNR evaluation across unsampled directions without increasing dataset resolution.

We note that the offline optimization runtime is dominated by NN inference and the number of SA iterations required for convergence. In our experiments, a single SA optimization required $\mathcal{O}(10^3)-\mathcal{O}(10^4)$ NN forward passes (corresponding to the $i_{\max} = 2000$ iterations used in Table 1), with the number of optimizations scaling linearly with the number of required beams and an additional step used to generate all nulls. On a standard CPU, this resulted in runtimes ranging from several seconds to a few minutes per pattern. In practical deployments, NN inference would be executed on lightweight embedded hardware or a small accelerator, reducing inference latency to the microsecond-millisecond range and enabling substantially faster offline optimization. The lookup table itself contained approximately 1-3 entries per scenario in our experiments, with each entry storing a $\boldsymbol{W}$, its associated beam and null directions, and the corresponding SLNR value. The table grows only when new UE directions or SLNR targets are introduced, and although larger RISs may require finer angular resolution due to narrower beamwidths, the growth remains linear in the number of desired beam directions rather than in the number of RIS elements.

IV. SIMULATIONS AND OBSERVATIONS

The RIS and environment models described in the previous sections were simulated using MATLAB under various scenarios to assess the ability of the proposed AI-driven framework in this work and in [7] to generalize for different channel conditions, RIS configurations, and dataset sizes. We consider the SLNR maximization problem for the $M = 100$ element RIS described earlier using the S/H implementation, where we attempt to create a single or multiple beams with either no nulls, a single null, or multiple nulls, with various training dataset sizes $\mathcal{D}$, numbers of BSW frequencies N, and channel models. We compare the results with the same cases discussed in [7] where the envelope detector circuit is used to sample the standing waves.

The RIS behaves as a single-row uniform linear array with $M = 100$ metallic patches periodically separated by $d_x = 20$ mm with $L_{\text{left}} = L_{\text{right}} = 0.5d_x$. The total length of the TL along x is therefore $L_{\text{tot}} = 100d_x$. To reduce the spatial sensitivity in realizing the biasing TL, while also reducing the fundamental frequency of the standing waves, we use

TABLE 2. NN architectures optimized by the GA.

Test Case	Number of Hidden Layers	Activation Functions	Number of Nodes per Layer	Number of Epochs	Batch Size	Validation Loss
Rectifier, $N = 25$, LoS, $D = 10^5$	5	{PReLU, tanh, PReLU, PReLU, ReLU}	{1024, 64, 512, 2048, 1024}	147	128	8.4779E-4
S/H, $N = 25$, LoS, $D = 10^5$	6	{PReLU, PReLU, PReLU, PReLU, tanh, PReLU}	{128, 512, 2048, 2048, 256, 512}	177	64	1.3813E-3
S/H, $N = 25$, LoS, $D = 10^4$	3	{ReLU, PReLU, PReLU}	{2048, 64, 512}	236	256	1.1148E-2
S/H, $N = 50$, LoS, $D = 10^5$	6	{PReLU, PReLU, PReLU, PReLU, tanh, PReLU}	{256, 2048, 2048, 2048, 2048, 2048, 1024}	134	256	1.4903E-3
S/H, $N = 25$, Rician $\kappa = 10$, $D = 10^4$	3	{ReLU, tanh, PReLU}	{2048, 256, 2048}	164	512	1.0347E-2
S/H, $N = 25$, Rician $\kappa = 5$, $D = 10^4$	3	{PReLU, tanh, PReLU}	{2048, 256, 2048}	164	512	9.1089E-3
S/H, $N = 25$, Rician $\kappa = 1$, $D = 10^4$	3	{ReLU, PReLU, ReLU}	{2048, 64, 2048}	198	128	7.0566E-3
S/H, $N = 25$, Rayleigh, $D = 10^4$	1	{ReLU}	{2048}	198	128	4.1178E-3

a serpentine structure for the TL as in Fig. 2, such that $L_p = 131.42$ mm. The TL is biased by a superposition of N standing waves with DC operating point $W_0 = 6.4$ V, around which point the reflection phase versus biasing voltage is approximately linear, as seen in Fig. 5(b). The standing-wave fundamental frequency is $f_b \approx 3.87$ MHz and the carrier frequency of the transmitted signal is $f_c = 2.45$ GHz. We set the transmit power $P = 0$ dBm and noise variance $\sigma_n^2 = 1$. We choose the sampling instant $t_0 = 8/\omega_b$ to satisfy $\sin(n\omega_b t_0) \neq 0$ for all $n \in \{1, 2, \ldots, N\}$ when $N \leq 50$. We assume that all S/H circuits are synchronized and sample the standing waves at the same time.

We first created the datasets for various scenarios and trained the NNs using the GA discussed previously. In Algorithm 1, we use $\sigma_1 = 0.008$ V, $\sigma_2 = 0.8$ V, $N \in \{25, 50\}$, and training dataset size $\mathcal{D} \in \{10^4, 10^5\}$ depending on the specific simulated scenario. The varactors may be biased between 4 V and 15 V, according to Fig. 5(a). The UE directions are distributed uniformly at azimuth angles between $-60°$ and $60°$ with respect to the broadside of the RIS, with a spacing of $1.5°$, which is less than the simulated 3 dB beamwidth of the given 100-element RIS configuration. Thus, $n_{\text{angles}} = 81$. Generally, the beamwidth would decrease and more sampled directions would be required to accurately capture beams and nulls as the size of the RIS increases. In such cases, the operator may choose to limit the data collection to only specific angle regions of interest to reduce the complexity of data collection, as well as training and modeling difficulties using the NN.

Other relevant parameters used in the GA and SA are summarized in Table 1. Note that the GA parameters are given as sets of hyperparameters, used to randomly initialize the architectures of the individual NNs in the populations. The NNs used in the GA are defined and trained using TensorFlow Keras [22] in Python, using the Adam Optimizer [23] and mean squared error as the loss function. The learning rate is reduced by half if the validation loss stops improving after 20 consecutive epochs and the training is terminated if the validation loss does not improve after 30 consecutive epochs. Thus, the effective number of epochs may be lower than the nominal number of epochs selected from Table 1. Lastly, the NNs have a regularization penalty of 5×10^{-7} to avoid overfitting [24]. The training sets are the first 90% of the datasets and the validation sets consist of the last 10%. The fitness of each MLP in the GA is determined by its validation loss.

The network architectures that result in the smallest validation losses after GA convergence are outlined in Table 2 and the following observations are made:

1) The optimal NNs contain mainly PReLU, tanh, and ReLU units, while the sigmoid activation function is in none of the optimal architectures [17]. The GA can potentially be made more efficient by excluding the sigmoid and ReLU activation functions to improve convergence and reduce the dimensionality of the optimization problem. The PReLU function exhibits better flexibility than ReLU since it enables gradients when the input values are negative, with a learned slope for $y = ax$ when $x < 0$, thus fixing the "dying neuron" problem that ReLU has [25], [26]. If the slope is very small, the negative values may be absorbed in the other NN layers and become negligible.
2) Validation losses are smallest for the largest training sets.
3) The deepest NNs correspond to the largest training sets, showcasing the complexity of the model required to provide a sufficiently generalized representation of the environment.
4) Validation losses remain relatively low for Rician and Rayleigh fading as the NLoS effects become more

TABLE 3. Average SLNR [dB] after SA convergence using the optimal NN models and the lookup table for various beam and null directions.

Beam 1	Beam 2	Null 1	Null 2	Rectifier, $N = 25$, LoS, $\mathcal{D} = 10^5$ [7]	S/H, $N = 25$, LoS, $\mathcal{D} = 10^5$	S/H, $N = 25$, LoS, $\mathcal{D} = 10^4$	S/H, $N = 50$, LoS, $\mathcal{D} = 10^5$	S/H, $N = 25$, Rician $\kappa = 10$, $\mathcal{D} = 10^4$	S/H, $N = 25$, Rician $\kappa = 5$, $\mathcal{D} = 10^4$	S/H, $N = 25$, Rician $\kappa = 1$, $\mathcal{D} = 10^4$	S/H, $N = 25$, Rayleigh, $\mathcal{D} = 10^4$
-40.5°				33.67	34.10	33.69	35.77	31.77	31.56	26.24	16.60
25.5°				33.92	35.12	32.49	36.51	30.65	31.79	27.23	16.52
-30°		-49.5°		27.68	33.00	21.45	35.60	9.50	20.95	14.13	-0.28
-19.5°	-49.5°			28.35	29.80	26.91	30.75	26.40	27.07	19.99	12.93
-19.5°	49.5°	10.5°	-40.5°	23.02	25.09	6.98	25.85	9.80	10.97	-3.05	-6.02
30°		0°		5.33	21.39	5.86	25.51	7.28	8.11	10.35	-0.26
-35°		-50°		31.41	32.19	16.34	31.46	9.69	14.09	9.34	-0.26
-20°	50°	10°	-40°	25.40	22.84	13.09	28.46	10.77	11.05	1.51	-6.07

significant, but this may be a consequence of the fact that the radiation pattern values approach an average value (similar to a noise floor), whereas the shapes of the beams or nulls become harder to distinguish since the beampattern is no longer predominantly characterized by them (see Fig. 7) [27], [28]. Since we define the NLoS channels as $\boldsymbol{h}_{\text{NLoS}} \in \mathcal{CN}(0, 1)^{M\times 1}$ in (6) and $M = 100$, this average value would generally be $P_{\text{avg}} \leq 10\log(M) = 20$ dB for the Rayleigh case. See the Appendix for derivation. In the results shown in Table 3, we see an average SNR value around 16.6 dB for both cases where the channels are characterized by Rayleigh fading and a single beam is to be formed, confirming this observation.

5) Each of the NNs has its own unique architecture and hyperparameters which lead to the lowest validation loss for each case. This further emphasizes the need for automation in designing the NNs by using the GA discussed in this paper.

After creating the optimal NNs, we used the SA algorithm along with the lookup table to optimize the beam patterns by creating beams and nulls at specific directions. The results are summarized in Table 3, where SA was executed 10 times for each test case and the SLNR [dB] shown for each case is obtained by averaging over 1000 different channel realizations.

We make the following remarks from the simulation results:

1) As expected, the S/H-based RIS achieves better beamforming or nullforming performance compared to the envelope detector implementation. For LoS channels with $\mathcal{D} = 10^5$, S/H achieves higher SLNR in all cases except the last one. Furthermore, as seen in Fig. 8, the S/H model is also able to eliminate the spurious broadside beam, thus minimizing interference from the BS direction.
2) With smaller datasets, the S/H-based model achieves beamforming performance comparable to the model trained on the larger datasets. The nullforming ability degrades due to the complexity of orthogonalizing the RIS response to the channels when limited data is available. However, the training and NN optimization times are much smaller with the reduced datasets.
3) The linear interpolation procedure does not apply for NLoS-dominant channels because the radiation pattern is no longer distinguished by well-characterized beams. This is especially evident in Fig. 7 where, as the random NLoS components become stronger, the radiation pattern exhibits more random behavior, until finally in the Rayleigh case the beams no longer exist. This point will be elaborated on further below.
4) Adding more BSW frequencies yields a marginal boost in SLNR. The system can be simplified by using fewer BSW frequencies, depending on the range of desired beamforming angles. This topic is further discussed in [3, Section V-E.5].

To better understand the effects seen in Fig. 7, we refer back to (6) where the channel models are defined. For a very high Rician factor $\kappa \to \infty$, the channels are dominated by the LoS components, resulting in a radiation pattern characterized primarily by beams and nulls [29]. The wave-controlled RIS can be used for beamforming or nullforming either analytically or using heuristic algorithms, as seen in [3] when the channels are well known, or using machine learning, as in this work and in [7]. When κ decreases, the scattered multipath (NLoS) components become more significant and the peaks created by the LoS channels become less prominent. This scattering also negatively affects the ability to form nulls, as demonstrated by the SLNR degradation seen in Table 3. Ultimately, for every small κ, the radiation pattern is no longer characterized by beams and nulls. The SLNR objective function serves

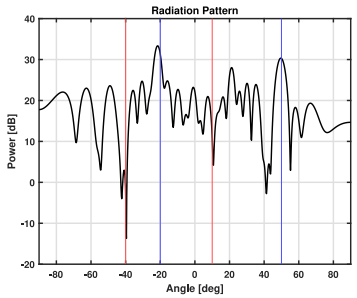

(a) LoS, $D = 10^5$, SLNR = 17.36 dB

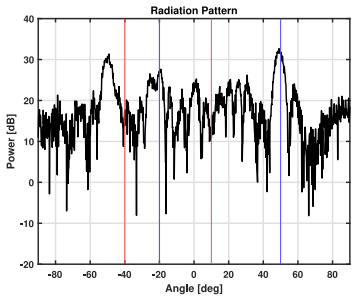

(b) Rician $\kappa = 10$, $D = 10^4$, SLNR = 8.63 dB

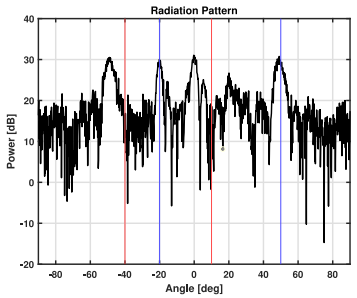

(c) Rician $\kappa = 5$, $D = 10^4$, SLNR = 9.25 dB

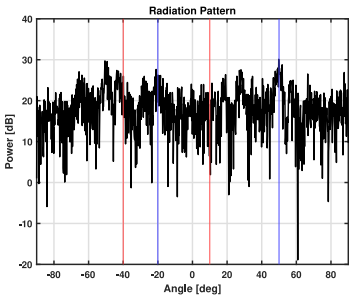

(d) Rician $\kappa = 1$, $D = 10^4$, SLNR = 1.07 dB

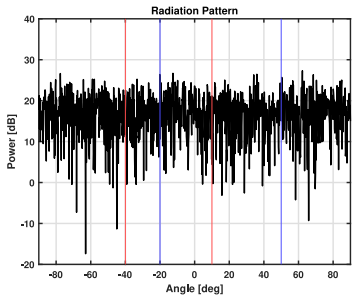

(e) Rayleigh, $D = 10^4$, SLNR = -6.82 dB

FIGURE 7. Converged radiation patterns using the S/H-based RIS, maximizing received powers at $-20°$ and $50°$ and minimizing received powers at $10°$ and $-40°$. $N = 25$ with various dataset sizes and channel models for h_{BS} and h_k.

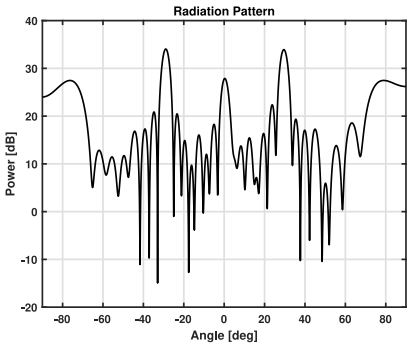

(a) Envelope detector, SLNR = 5.96 dB [7]

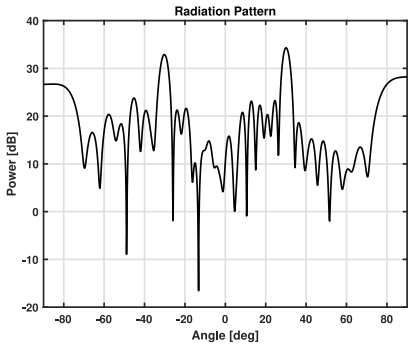

(b) S/H, SLNR = 24.46 dB

FIGURE 8. Converged radiation patterns, attempting to create a beam at $30°$ and a null at $0°$. $N = 25$, $D = 10^5$, LoS channels.

to maximize the received power towards desired directions by configuring the RIS to constructively add the signals at the UE and create a coherent sum, while minimizing the power for other directions by orthogonalizing the RIS response to the cascaded channel response. These properties are not visible in the beampattern. The extreme case of Rayleigh fading, where there is no LoS component, is shown in Fig. 7(e).

An observed strength of the AI-driven approach in this work is the resistance of the NN to random NLoS components when modeling the radiation patterns. Although the patterns contain stronger NLoS components as κ decreases in Fig. 7, there are still distinguishable peaks in the received power at directions of interest due to the LoS components. Thus, the NN is able to learn the LoS components of the radiation pattern while filtering out the randomness caused by the NLoS components, proving its resilience to channel changes while the beam directions remain constant. However, the SLNR value degrades significantly when attempting to also generate nulls in certain directions, due to the sensitivity of forming sharp nulls towards the LoS components and the difficulty of orthogonalizing the RIS response to the NLoS components. When the LoS paths are stronger, it becomes easier to use the RIS to maximize the difference in power between the desired beams and nulls.

In the Rayleigh case, the NN is unable to differentiate between the true and estimated radiation patterns since the zero-mean random channel coefficients encourage the NN to estimate the mean value of the radiation pattern in order to minimize the MSE. Thus, the relationship between the input BSWs and the output radiation pattern is not learned, and the AI-driven approach fails. Therefore, to

maximize the signal power towards particular directions in the Rayleigh case, codebook approaches may be used, such as those described in [30], [31], [32], [33]. Alternatively, if an accurate system model is available and the channels can be estimated, the algorithms presented in [3] and below can be utilized. However, forming nulls for channels with significant NLoS components is generally not possible, unless feedback is available from a cooperative receiver at which low interference is desired.

V. OPTIMIZATION UNDER IDEAL CONDITIONS

Until now, we have assumed that the RIS model does not admit a closed-form characterization, CSI of the flat-fading wireless channels is unavailable, and there is no live feedback from the user-end receivers. We have seen the capabilities and limitations of the AI-based approach in modeling the entire system using an NN alongside various heuristic algorithms to optimize SLNR. We now compare this approach with a more idealized case, where the mathematical model of the RIS is known and can be used for optimization, and CSI for the cascaded BS-RIS and each RIS-UE channel is perfectly estimated. We use the "Combined Algorithm" presented in [3] to optimize the SLNR metric for the same S/H-based RIS with $N = 25$ BSWs and the channels presented in Table 3. The Combined Algorithm is outlined as follows:

1) *Ideal RIS reflection phase optimization* – For simple SNR maximization problems, the reflection phases are calculated as the conjugates of the cascaded channels to produce a coherent sum. For channels $h_{\text{BS}}(m) = a(m)e^{j\alpha(m)}$ and $h_k = b_k(m)e^{j\beta_k(m)}$, the reflection phase that maximizes the power is $\angle\Gamma = e^{-j(\alpha(m)+\beta_k(m))}$ under the assumption that $|\Gamma| = 1$, yielding

$$\begin{aligned}\text{SNR}_k &= \frac{P\left|\boldsymbol{h}_{\text{BS}}^T\boldsymbol{\Gamma}\boldsymbol{h}_k\right|^2}{\sigma_n^2} \\ &= \frac{P\left|\sum_{m=0}^{M-1} h_{\text{BS}}(m)\Gamma(m)h_k(m)\}\right|^2}{\sigma_n^2} \\ &= \frac{P\left|\sum_{m=0}^{M-1} a(m)b_k(m)\right|^2}{\sigma_n^2}. \end{aligned} \tag{15}$$

When there are multiple desired UE receivers, the above procedure is repeated for every receiver k, creating reflection coefficients denoted by $\boldsymbol{\Gamma}_k$ and averaged as $\boldsymbol{\Gamma}_{\text{avg}} = \frac{1}{K}\sum_{k=0}^{K-1}\boldsymbol{\Gamma}_k$. To enforce the unity gain condition, the resulting reflection coefficient is given by $\boldsymbol{\Gamma} = e^{j\angle\boldsymbol{\Gamma}_{\text{avg}}}$.

2) *Creating beampattern nulls* – For the j-th desired null direction, the cascaded channel is calculated as $r_j(m) = h_{\text{BS}}(m)\Gamma(m)h_j(m)$. We subtract each $r_j(m)$ value by the average $\bar{r}_j = \frac{1}{M}\sum_{m=0}^{M-1} r_j(m)$ to obtain the reflection phases that collectively orthogonalize the channel response

$$\Gamma(m) = \exp\left(j\angle\left(\frac{r_j(m) - \bar{r}_j}{h_{\text{BS}}(m)h_j(m)}\right)\right). \tag{16}$$

The process is repeated for every null direction j until the largest $\bar{r}_j$ is below a desired threshold.

3) *Phase to voltage conversion* – The reflection phases $\angle\Gamma(m)$ are mapped to the varactor reverse bias voltages $v(m)$ using the one-to-one function in Fig. 5(b). Reflection phases outside of the 4-15 V range are clipped.

4) *Voltage to BSW amplitude mapping using Weighted Least Squares (WLS)* – Linear regression is performed to obtain BSW amplitudes $\boldsymbol{W}$ that would produce a standing wave with samples $v(m)$ for $m = 0, 1, \ldots, M-1$. This is done by solving the convex optimization problem

$$\min_{\boldsymbol{W}} \sum_{m=0}^{M-1} \gamma(m)||w(m) - v(m)||_2^2, \tag{17}$$

where $\gamma(m)$ is a weighting factor based on the rate of change of the RIS phase with respect to $v(m)$, and $w(m)$ is the resulting DC voltage at each varactor location m when the BSWs on the transmission line are sampled by the S/H.

5) *SLNR maximization using Simulated Annealing* – The SA algorithm from before is used to maximize the SLNR by randomly perturbing $\boldsymbol{W}$ and calculating the resulting SLNR with the given channel coefficients. This is a crucial step when designing the nulls in the radiation pattern.

The results of the above approach for the same scenarios as in Table 3 and Fig. 7 are given in Table 4, also averaged over 1000 channel realizations, and Fig. 9. Overall, the wave-controlled RIS effectively enhances the SLNR for all cases shown, for both the desired beam and null directions. A degradation of approximately only 4-6 dB is seen between pure LoS and Rayleigh channels, underlining the robustness of the Combined Algorithm to the channel model. We also note the SLNR improvement of this baseline compared to the data-driven approach shown in Table 3, specifically in the cases where nullforming is involved, showcasing the delicacy required in system modeling to enhance the SLNR objective.

We conclude that the main performance bottleneck in the optimization is the accurate modeling of the RIS and environment (including channel coefficients). Both the idealized and AI-based approaches achieve beamforming and power gains using the same SA algorithm. However, the nullforming capability of the AI framework degrades significantly for channel models with NLoS randomness due to the sensitivity of the nulls to small changes in RIS element biasing, requiring careful modeling. This is often impractical when using an NN trained on a limited dataset when the

TABLE 4. Average SLNR [dB] after convergence of the Combined Algorithm with perfect CSI for various channel models, using the S/H-based RIS with $N = 25$ BSWs.

Beam 1	Beam 2	Null 1	Null 2	LoS	$\kappa = 10$	$\kappa = 5$	$\kappa = 1$	$\kappa = 0$
-40.5°				34.42	34.27	34.00	32.48	30.96
25.5°				35.52	35.23	34.85	32.92	30.99
-30°		-49.5°		34.57	33.85	33.34	30.96	28.67
-19.5°	-49.5°			30.73	30.63	30.49	29.27	27.96
-19.5°	49.5°	10.5°	-40.5°	29.13	28.41	28.04	25.25	23.62
30°		0°		34.50	33.93	33.39	30.97	28.72
-35°		-50°		34.28	33.66	33.09	30.65	28.61
-20°	50°	10°	-40°	28.30	28.14	27.75	25.04	23.61

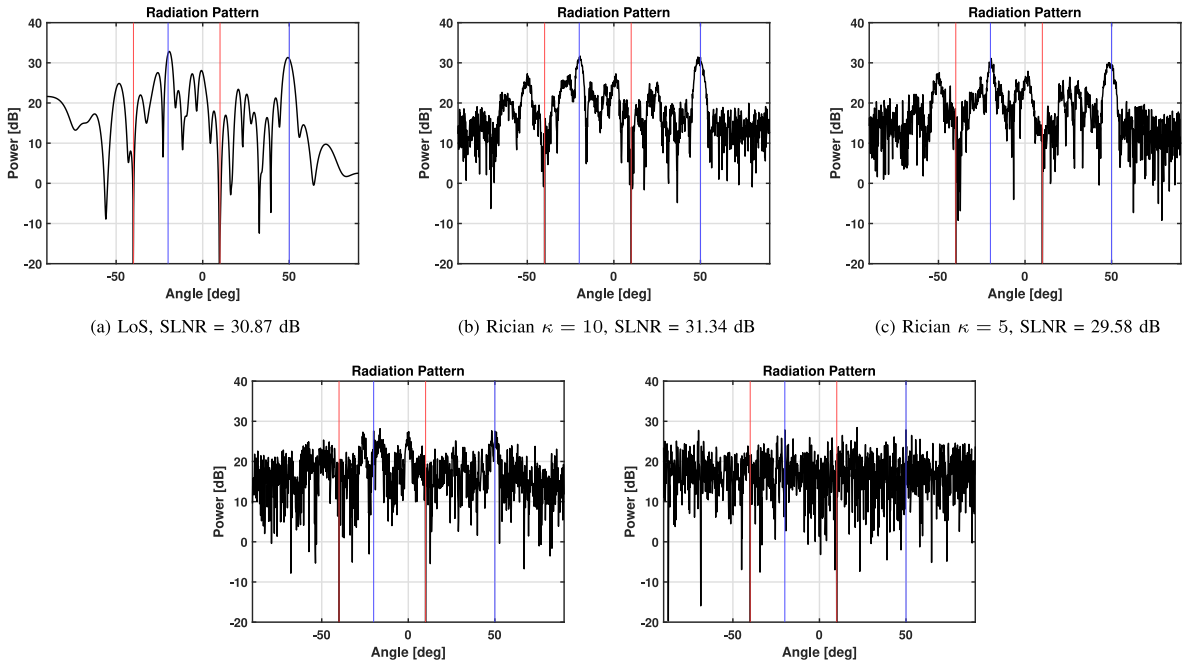

(a) LoS, SLNR = 30.87 dB (b) Rician $\kappa = 10$, SLNR = 31.34 dB (c) Rician $\kappa = 5$, SLNR = 29.58 dB

(d) Rician $\kappa = 1$, SLNR = 27.53 dB (e) Rayleigh, SLNR = 27.86 dB

FIGURE 9. Converged radiation patterns using the S/H-based RIS, using the Combined Algorithm with perfect CSI, maximizing received powers at $-20°$ and $50°$ and minimizing received powers at $10°$ and $-40°$. $N = 25$ with various channel models for h_{BS} and h_k.

mapping between a large input space (N-dimensional) and a much larger output space (number of sampled directions) are related via nonlinear and complex functions. Although the proposed NN demonstrated the ability to model an entire radiation pattern for beamforming, a smaller portion of the radiation pattern can be used for SLNR maximization in cases where only certain angle regions are of interest, to limit the size of the model and the need for excessively large datasets or NN architectures.

VI. CONCLUSION

This work has demonstrated the practicality and robustness of the data-driven optimization framework proposed in [7] for diverse RIS implementations and channel models. Through comparative analysis, we showed that the S/H-based wave-controlled RIS provides superior control over radiation patterns compared to the rectifier-based RIS. Specifically, it achieves higher SLNR values under identical constraints and effectively eliminates the spurious broadside beam in LoS channels, underscoring its enhanced beamforming capability.

We further investigated the impact of dataset size on performance and found that reducing the training data does not significantly impair the RIS's modeling ability. However, nullforming performance in LoS channels, particularly in configurations involving a 100-element RIS, exhibits sensitivity to dataset size due to the precision required in modeling deep nulls. This suggests that more extensive training data may be necessary for accurate NN modeling in such cases. Nevertheless, the reduction in dataset size yields substantial improvements in training and optimization

efficiency, thereby enhancing the overall feasibility of the proposed approach.

The adaptability of the optimization algorithms to varying channel conditions was also examined. Despite the absence of channel state information, the algorithms maintain consistent beamforming performance for Rician and LoS channels. The NN was able to predict the LoS components of the channel as functions of the RIS biasing, allowing better control over the radiation patterns and approaching the performance of ideal cases. In Rayleigh and severe NLoS fading environments, due to increased multipath scattering and the absence of a dominant LoS path, maximizing and minimizing received powers at specific directions becomes more difficult when using the NN to model the entire radiation pattern. Although the AI-driven approach has more difficulty in systematically modeling cases where NLoS multipath is dominant, the wave-controlled RIS in general was able to maximize the SLNR metric under ideal conditions where channel coefficients and a precise RIS model are available for all cases.

APPENDIX
AVERAGE RECEIVED POWER USING RIS UNDER RAYLEIGH FADING

From (9), we define the received signal power at receiver k as

$$P_k = P|\boldsymbol{h}_{\text{BS}}^T \boldsymbol{\Gamma} \boldsymbol{h}_k|^2, \tag{18}$$

which is equivalent to

$$\begin{aligned} P_k &= \left(\sum_{i=1}^{M} \Gamma_i h_{\text{BS},i} h_{k,i}\right)\left(\sum_{i=1}^{M} \Gamma_i^* h_{\text{BS},i}^* h_{k,i}^*\right) \\ &= \left(\sum_{i=1}^{M}\sum_{j=1}^{M} \Gamma_i \Gamma_j^* h_{\text{BS},i} h_{\text{BS},j}^* h_{k,i} h_{k,j}^*\right), \end{aligned} \tag{19}$$

where M is the number of RIS elements, $\boldsymbol{\Gamma}$ is the $M \times M$ diagonal matrix containing the RIS reflection coefficients, $\boldsymbol{h}_{\text{BS}} \in \mathcal{CN}(0, \sigma_{\text{BS}}^2)^{M\times 1}$, $\boldsymbol{h}_k \in \mathcal{CN}(0, \sigma_k^2)^{M\times 1}$, and P is the transmit power. Note that $E[h_{\text{BS},i} h_{\text{BS},j}^*] = \sigma_{\text{BS}}^2 \delta_{ij}$ and $E[h_{k,i} h_{k,j}^*] = \sigma_k^2 \delta_{ij}$ since both $\boldsymbol{h}_{\text{BS}}$ and $\boldsymbol{h}_k$ are independent and identically distributed (i.i.d.) circularly symmetric complex Gaussian random variables, and $E[h_{\text{BS},i} h_{k,i}] = 0$ since the two distributions are independent.

The received power is therefore given by

$$P_k = P\sum_{i=1}^{M} |\Gamma_i|^2 |h_{\text{BS},i}|^2 |h_{k,i}|^2. \tag{20}$$

To find the average power, we take the expectation of (20) which yields

$$\begin{aligned} E[P_k] &= E\left[P\sum_{i=1}^{M} |\Gamma_i|^2 |h_{\text{BS},i}|^2 |h_{k,i}|^2\right] \\ &= P\sum_{i=1}^{M} E\left[|\Gamma_i|^2\right] E\left[|h_{\text{BS},i}|^2\right] E\left[|h_{k,i}|^2\right], \end{aligned} \tag{21}$$

where $E[|h_{\text{BS},i}|^2] = \sigma_{\text{BS}}^2$, $E[|h_{k,i}|^2] = \sigma_k^2$, and $E[|\Gamma_i|^2] = |\Gamma_i|^2$. Since $|\Gamma_i|^2 \leq 1$, and for the case where $\sigma_{\text{BS}}^2 = \sigma_k^2 = 1$ and the transmit power $P = 1$, we can express the average received power as

$$E[P_k] = P\sum_{i=1}^{M} |\Gamma_i|^2 \sigma_k^2 \sigma_{\text{BS}}^2 \leq M. \tag{22}$$

ACKNOWLEDGMENT

The authors thank DS SIMULIA for providing CST Studio Suite.

REFERENCES

[1] C. Pan et al., "Reconfigurable intelligent surfaces for 6G systems: Principles, applications, and research directions," *IEEE Commun. Mag.*, vol. 59, no. 6, pp. 14–20, Jun. 2021.
[2] E. Ayanoglu, F. Capolino, and A. Swindlehurst, "Wave-controlled metasurface-based reconfigurable intelligent surfaces," *IEEE Wireless Commun.*, vol. 29, no. 4, pp. 86–92, Aug. 2022.
[3] G. Ben-Itzhak, M. Saavedra-Melo, B. Bradshaw, E. Ayanoglu, F. Capolino, and A. Swindlehurst, "Design and operation principles of a wave-controlled reconfigurable intelligent surface," *IEEE Open J. Commun. Soc.*, vol. 5, pp. 7730–7751, 2024.
[4] M. Saavedra-Melo, B. Bradshaw, V. Yao, E. Ayanoglu, A. Swindlehurst, and F. Capolino, "Experimental analysis of biasing voltage generation in wave-controlled RIS," 2025, *arXiv:2509.07293*.
[5] A. Swindlehurst, G. Zhou, R. Liu, C. Pan, and M. Li, "Channel estimation with reconfigurable intelligent surfaces—A general framework," *Proc. IEEE*, vol. 110, no. 9, pp. 1312–1338, Sep. 2022.
[6] C. Pan et al., "An overview of signal processing techniques for RIS/IRS-aided wireless systems," *IEEE Sens. J. Selected Topics Signal Process.*, vol. 16, no. 5, pp. 883–917, Aug. 2022.
[7] G. Ben-Itzhak, M. Saavedra-Melo, E. Ayanoglu, F. Capolino, and A. Swindlehurst, "AI-driven optimization of wave-controlled reconfigurable intelligent surfaces," *IEEE Open J. Commun. Soc.*, vol. 6, pp. 6650–6665, 2025.
[8] B. Saglam, D. Gurgunoglu, and S. S. Kozat, "Deep reinforcement learning based joint downlink beamforming and RIS configuration in RIS-aided MU-MISO systems under hardware impairments and imperfect CSI," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, 2023, pp. 66–72.
[9] C. Huang, R. Mo, and C. Yuen, "Reconfigurable intelligent surface assisted multiuser MISO systems exploiting deep reinforcement learning," *IEEE Sens. J. Sel. Areas Commun.*, vol. 38, no. 8, pp. 1839–1850, Aug. 2020.
[10] W. Wang and W. Zhang, "Intelligent reflecting surface configurations for smart radio using deep reinforcement learning," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 8, pp. 2335–2346, Aug. 2022.
[11] P. Horowitz and W. Hill, *The Art of Electronics*. Cambridge, U.K.: Cambridge Univ., 1980.
[12] F. Costa and M. Borgese, "Electromagnetic model of reflective intelligent surfaces," *IEEE Open J. Commun. Soc.*, vol. 2, pp. 1577–1589, 2021.
[13] D. Hanna, M. S. Melo, F. Shan, and F. Capolino, "A versatile polynomial model for reflection by a reflective intelligent surface with varactors," in *Proc. IEEE Int. Symp. Antennas Propag. USNC-URSI Radio Sci. Meeting (AP-S/URSI)*, 2022, pp. 679–680.
[14] F. Capolino, A. Vallecchi, and M. Albani, "Equivalent transmission line model with a lumped X-circuit for a metalayer made of pairs of planar conductors," *IEEE Trans. Antennas Propag.*, vol. 61, no. 2, pp. 852–861, Feb. 2013.
[15] Q. Wu and R. Zhang, "Beamforming optimization for intelligent reflecting surface with discrete phase shifts," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2019, pp. 7830–7833.
[16] E. Basar, M. Di Renzo, J. D. M. Debbah, M.-S. Alouini, and R. Zhang, "Wireless communications through reconfigurable intelligent surfaces," *IEEE Access*, vol. 7, pp. 116753–116773, 2019.
[17] B. Ding, H. Qian, and J. Zhou, "Activation functions and their characteristics in deep neural networks," in *Proc. Chinese Control Decision Conf. (CCDC)*, 2018, pp. 1836–1841.

[18] L. N. Smith, "A disciplined approach to neural network hyper-parameters: Part 1—Learning rate, batch size, momentum, and weight decay," 2018, *arXiv:1803.09820*.
[19] S. Forrest, "Genetic algorithms," *ACM Comput. Surv.*, vol. 28, no. 1, pp. 77–80, 1996.
[20] J. V. Domashova, S. S. Emtseva, V. S. Fail, and A. S. Gridin, "Selecting an optimal architecture of neural network using genetic algorithm," *Procedia Comput. Sci.*, vol. 190, 2021, pp. 263–273.
[21] V. Cerny, "Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm," *J. Optim. Theory Appl.*, vol. 45, pp. 41–51, Jan. 1985.
[22] M. Abadi et al., "TensorFlow: Large-scale machine learning on heterogeneous systems," in *Proc. 12th USENIX Conf. Oper. Syst. Design Implement.*, 2015, pp. 265–283. [Online]. Available: https://www.tensorflow.org/
[23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017, *arXiv:1412.6980*.
[24] G. Shi, J. Zhang, H. Li, and C. Wang, "Enhance the performance of deep neural networks via L2 regularization on the input of activations," *Neural Process. Lett.*, vol. 50, no. 1, pp. 57–75, Aug. 2019.
[25] Y. Bai, "ReLU-function and derived function review," in *Proc. SHS Web Conf.*, vol. 144, 2022, p. 2006.
[26] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1026–1034.
[27] F. Digham, M.-S. Alouini, and M. Simon, "On the energy detection of unknown signals over fading channels," in *Proc. IEEE Int. Conf. Commun.*, vol. 5, 2003, pp. 3575–3579.
[28] B. Sklar, "Rayleigh fading channels in mobile digital communication systems. I. characterization," *IEEE Commun. Mag.*, vol. 35, no. 7, pp. 90–100, Sep. 1997.
[29] B. Van Veen and K. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE ASSP Mag.*, vol. 5, no. 2, pp. 4–24, Apr. 1988.
[30] 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.
[31] J. An et al., "Codebook-based solutions for reconfigurable intelligent surfaces and their open challenges," *IEEE Wireless Commun.*, vol. 31, no. 2, pp. 134–141, Apr. 2024.
[32] Z. Yu, J. An, L. Gan, H. Li, and S. Chatzinotas, "Weighted codebook scheme for RIS-assisted point-to-point MIMO communications," *IEEE Wireless Commun. Lett.*, vol. 14, no. 5, pp. 1571–1575, May 2025.
[33] J. An, C. Xu, L. Gan, and L. Hanzo, "Low-complexity channel estimation and passive beamforming for RIS-assisted MIMO systems relying on discrete phase shifts," *IEEE Trans. Commun.*, vol. 70, no. 2, pp. 1245–1260, Feb. 2022.

GAL BEN-ITZHAK (Graduate Student Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from the University of California at Irvine, Irvine, CA, USA, in 2023 and 2024, respectively, where he is currently pursuing the Ph.D. degree. His current research and professional interests include high-speed communications, optimization, digital signal processing, high-speed circuit design, and signal integrity.

ENDER AYANOGLU (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1986. He was with the Bell Laboratories Communications Systems Research Laboratory, Holmdel, NJ, USA. From 1999 to 2002, he was a Systems Architect with Cisco Systems Inc., San Jose, CA, USA. Since 2002, he has been a Professor with the Department of Electrical Engineering and Computer Science, University of California at Irvine, Irvine, CA, USA, where he was the Director of the Center for Pervasive Communications and Computing and the Conexant-Broadcom Endowed Chair from 2002 to 2010. He was a recipient of the IEEE Communications Society Stephen O. Rice Prize Paper Award in 1995, the IEEE Communications Society Best Tutorial Paper Award in 1997, and the IEEE Communications Society Communication Theory Technical Committee Outstanding Service Award in 2014. In 2023, he received the IEEE Communications Society Joseph L. LoCicero Award for outstanding contributions to IEEE Communications Society journals as an Editor, the Editor-in-Chief (EiC), and the Founding EiC. From 2000 to 2001, he was the Founding Chair of the IEEE-ISTO Broadband Wireless Internet Forum, an industry standards organization. He served on the Executive Committee for the IEEE Communications Society Communication Theory Committee from 1990 to 2002, and its Chair from 1999 to 2002. From 1993 to 2014, he was an Editor of IEEE TRANSACTIONS ON COMMUNICATIONS. He was the Editor-in-Chief of IEEE TRANSACTIONS ON COMMUNICATIONS from 2004 to 2008 and the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS Series on Green Communications and Networking from 2014 to 2016. He was the Founding Editor-in-Chief of IEEE TRANSACTIONS ON GREEN COMMUNICATIONS AND NETWORKING from 2016 to 2020. He served as an IEEE Communications Society Distinguished Lecturer from 2022 to 2023 and from 2024 to 2025.

FILIPPO CAPOLINO (Fellow, IEEE) received the Laurea (cum laude) and the Ph.D. degrees in electrical engineering from the University of Florence, Italy, in 1993 and 1997, respectively. From 1997 to 1999, he was a Fulbright and then a Postdoctoral Fellow with the Department of Aerospace and Mechanical Engineering, Boston University, MA, USA. From 2000 to 2001, part of 2005 and in 2006, he was a Research Assistant Visiting Professor with the Department of Electrical and Computer Engineering, University of Houston, Houston, TX, USA. From 2002 to 2008, he was an Assistant Professor with the Department of Information Engineering, University of Siena, Italy. He has been a Visiting Professor with the Fresnel Institute, Marseille, France, 2003, and with the Centre de Recherche Paul Pascal, Bordeaux, France, 2010. He is currently a Professor with the Department of Electrical Engineering and Computer Science, University of California at Irvine, Irvine, CA, USA. In 2022, he held the title of Cathedra of Excellence with the University of Carlos III, Madrid, Spain. His current research interests include metamaterials and their applications, antennas and wireless systems, sensors in both microwave and optical ranges, millimeter wave technology, electron beam devices, plasmonics, microscopy, optical devices, and applied electromagnetics in general. He received the R. W. P. King Prize Paper Award from the IEEE Antennas and Propagation Society for the Best Paper of the Year 2000, by an author under 36. He is a Fellow of OPTICA. He was the Founder and EU Coordinator of the EU Doctoral Programs on Metamaterials from 2004 to 2009. From 2002 to 2008, he served as an Associate Editor for the IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION. He is an Editor of the *Metamaterials Handbook* (Boca Raton, FL, USA: CRC Press, 2009).

A. LEE SWINDLEHURST (Life Fellow, IEEE) received the B.S. and M.S. degrees in electrical engineering from Brigham Young University (BYU) in 1985 and 1986, respectively, and the Ph.D. degree in electrical engineering from Stanford University in 1991. He was with the Department of Electrical and Computer Engineering, BYU from 1990 to 2007, where he served as the Department Chair from 2003 to 2006. From 1996 to 1997, he held a joint appointment as a Visiting Scholar with Uppsala University and the Royal Institute of Technology, Sweden. From 2006 to 2007, he was on leave working as the Vice President of Research with ArrayComm LLC, San Jose, CA, USA. Since 2007, he has been with the Electrical Engineering and Computer Science Department, University of California Irvine, where he is a Distinguished Professor and currently serving as the Department Chair. His research focuses on array signal processing for radar, wireless communications, and biomedical applications. He received the 2000 IEEE W. R. G. Baker Prize Paper Award, the 2006 IEEE Communications Society Stephen O. Rice Prize in the Field of Communication Theory, the 2006, 2010, and 2021 IEEE Signal Processing Society's Best Paper Awards, the 2017 IEEE Signal Processing Society Donald G. Fink Overview Paper Award, a Best Paper award at the 2020 and 2024 IEEE International Conferences on Communications, the 2022 Claude Shannon-Harry Nyquist Technical Achievement Award from the IEEE Signal Processing Society, and the 2024 Fred W. Ellersick Prize from the IEEE Communications Society. From 2014 to 2017, he was also a Hans Fischer Senior Fellow in the Institute for Advanced Studies at the Technical University of Munich, and in 2016, he was elected as a Foreign Member of the Royal Swedish Academy of Engineering Sciences (IVA).