

AirCNN via Reconfigurable Intelligent Surfaces: Architecture Design and Implementation

Meng Hua, *Senior Member, IEEE*, Haotian Wu, *Member, IEEE*, and Deniz Gündüz, *Fellow, IEEE*

Abstract—This paper introduces AirCNN, a novel paradigm for implementing convolutional neural networks (CNNs) via over-the-air (OTA) analog computation. By leveraging multiple reconfigurable intelligent surfaces (RISs) and transceiver designs, we engineer the ambient wireless propagation environment to emulate the operations of a CNN layer. To comprehensively evaluate AirCNN, we consider two types of CNNs, namely classic two-dimensional (2D) convolution (Conv2d) and light-weight convolution, i.e., depthwise separable convolution (ConvSD). For Conv2d realization via OTA computation, we propose and analyze two RIS-aided transmission architectures: multiple-input multiple-output (MIMO) and multiple-input single-output (MISO), balancing transmission overhead and emulation performance. We jointly optimize all parameters, including the transmitter precoder, receiver combiner, and RIS phase shifts, under practical constraints such as transmit power budget and unit-modulus phase shift requirements. We further extend the framework to ConvSD, which requires distinct transmission strategies for depthwise and pointwise convolutions. Simulation results demonstrate that the proposed AirCNN architectures can achieve satisfactory classification performance. Notably, Conv2d MISO consistently outperforms Conv2d MIMO across various settings, while for ConvSD, MISO is superior only under poor channel conditions. Moreover, employing multiple RISs significantly enhances performance compared to a single RIS, especially in line-of-sight (LoS)-dominated wireless environments.

Index Terms—Over-the-air computation, multiple-input and multiple-output (MIMO), reconfigurable intelligent surface (RIS), convolutional neural network (CNN).

I. INTRODUCTION

Reconfigurable intelligent surfaces (RISs), also known as intelligent reflecting surfaces (IRSs), have emerged as a key enabling technology for a wide range of applications in 6G wireless communication systems [1]. Generally speaking, an RIS is a planar metasurface composed of numerous low-cost passive reflecting elements, such as varactor diodes. A key feature of RIS technology is its ability to dynamically adjust the amplitudes and/or phase shifts of incident signals in real-time, thereby reflecting them in a controlled manner. Representative studies [2]–[4] have demonstrated that RISs can significantly enhance network throughput and reduce overall network energy consumption by optimizing the phase shifts.

Beyond wireless communication, RISs have also been proposed for analog computation [5]–[7]. By adjusting the RIS phase shifts, an edge server can minimize model parameter aggregation errors. Furthermore, the massive array of reflecting

elements within an RIS can be regarded as trainable neurons in the layer of a neural network. Thus, RIS technology holds great promise for the realization of physical neural networks, offering faster computation and lower latency. Despite this potential, research on RIS-assisted physical neural networks remains limited [8]–[13]. Early work [8] and [9] developed multi-layer RIS structures devised for controlled laboratory environments for performing deep learning tasks. Work in [10] and [11] proposed RISs that enable edge inference by modeling the RIS-programmable wireless channel as hidden over-the-air (OTA) artificial neural network layers. In [12] and [13], the authors utilized RISs to program channel impulse responses, achieving one-dimensional (1D) and two-dimensional (2D) convolutional neural networks (CNNs), respectively. However, these works have not fully exploited the spatial degrees of freedom (DoFs) available for neural network emulation.

In this paper, we propose AirCNN, a novel OTA computation framework that emulates 2D CNNs through joint optimization of multi-RIS phase shifts, transmitter precoders, and receiver combiners. Unlike previous works [8]–[13], existing methods cannot be directly extended to realize 2D CNNs. Emulating 2D CNNs poses additional challenges, including incompatibility with traditional data architectures, transceiver design constraints, and transmission protocol limitations. To address these challenges, we investigate two types of 2D CNNs: the classic 2D convolution (Conv2d) and depthwise separable convolution (ConvSD). We propose corresponding transmission architectures and protocols for both RIS-assisted multiple-input single-output (MISO) and RIS-assisted multiple-input multiple-output (MIMO) systems. A comprehensive comparison between these two RIS-aided systems for physical CNN implementation is conducted, focusing on both performance and implementation overhead. Simulation results demonstrate that the proposed methods achieve satisfactory classification accuracy. Moreover, it is shown that multi-RIS setups significantly outperform single-RIS configurations, particularly in line-of-sight (LoS)-dominated wireless channels.

II. SYSTEM MODEL

Fig. 1(a) depicts a conventional CNN architecture, which comprises three modules: initial layers, a middle layer, and final layers. The initial and final layers may contain a large number of neural network layers depending on the specific application, while the middle layer corresponds to the CNN. This paper aims to emulate the CNN layer using wireless hardware to achieve a similar function as shown in

This work was supported by the SNS JU Project 6G-GOALS under the EU's Horizon Program with Grant 101139232.

The authors are with the Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, U.K. (e-mail: {m.hua,haotian.wu17,d.gunduz}@imperial.ac.uk).




Fig. 1: Illustration of conventional and OTA CNN architectures.




Fig. 2: A toy example of transforming a convolutional operation to a matrix multiplication operation.

Fig. 1(b). In the analog-based CNN architecture, the middle layer consists of precoders at the transmitter, multiple RISs deployed over the air, and combiners at the receiver. We assume that the transmitter and receiver are equipped with N_t and N_r antennas, respectively. Furthermore, L RISs are deployed, each comprising M/L reflecting elements, where M is the total number of reflecting elements. Let $\bar{\mathbf{H}}_i$ and $\hat{\mathbf{H}}_i$ denote the complex equivalent baseband channel matrices from the transmitter to RIS i and from RIS i to the receiver, respectively. The phase-shift matrix of RIS i is denoted by $\Theta_i = \text{diag}(e^{j\theta_{i,1}}, e^{j\theta_{i,2}}, \dots, e^{j\theta_{i,M/L}})$, with $\theta_{i,m}$ denoting the m -th phase shift.

According to the definition of convolution, it is impossible to directly implement CNNs OTA without transformation. Unlike convolution, matrix multiplication can be inherently realized via OTA transmission. Therefore, we transform the convolutional operation into a multiplication operation by rearranging the matrices. To clearly illustrate this concept, Fig. 2 presents a toy example in which $\mathbf{X}$ and $\mathbf{W}$ represent a 3×3 input image matrix and a 2×2 convolutional kernel matrix, respectively. The top part of Fig. 2 shows the standard convolutional operation with a stride of 1 and no padding, while the bottom part demonstrates the equivalent multiplication matrix operation. Specifically, matrices $\mathbf{W}$ and $\mathbf{Y}$ are unfolded into vectors $\bar{\mathbf{W}}$ and $\bar{\mathbf{Y}}$, respectively, and rearranged through a piece-wise block vectorization approach. To accommodate practical hardware constraints, the numbers of transmit and receive antennas are set to 4 and 1, respectively, matching the unfolded matrix size of 4×1 . Thus, after data rearrangement, the convolutional operation can be physically

realized. Although Fig. 2 illustrates the case of a single kernel, multiple kernels can be extended similarly.

A key challenge in AirCNN is the joint design of the transmit precoder, receive combiner, and RIS phase-shift matrices to accurately emulate a given digital convolution kernel $\bar{\mathbf{W}}$. Here, we consider a general case where multiple kernels are considered. Thus, $\bar{\mathbf{W}}$ is a matrix, where each row represents one kernel. This imitation problem can be formulated as

$$\min_{\mathbf{F}_1, \mathbf{F}_2, \Theta} \|\mathbf{F}_2 \mathbf{H} \mathbf{F}_1 - \bar{\mathbf{W}}\|_F^2 + \mathbb{E}_{\mathbf{n}} \left\{ \|\mathbf{F}_2 \mathbf{n}\|^2 \right\} \quad (1a)$$

$$\text{s.t. } \|\mathbf{F}_1\|_F^2 \leq P_{\max}, \quad (1b)$$

$$|\Theta_{i,i}| = 1, \quad i = 1, \dots, M/L, \quad (1c)$$

where $\mathbf{F}_1$ and $\mathbf{F}_2$ denote the transmit precoder and receive combiner, respectively, $P_{\max}$ is the transmit power budget, and $\mathbf{n}$ denotes additive white Gaussian noise, following $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I})$. The end-to-end channel matrix $\mathbf{H}$ is modeled as $\mathbf{H} = \sum_{i=1}^L \hat{\mathbf{H}}_i \Theta_i \bar{\mathbf{H}}_i$. Instead of solving problem (1) via conventional convex optimization techniques, the approach proposed in this paper is to optimize $\mathbf{F}_1$, $\mathbf{F}_2$, and Θ through end-to-end training based on the given loss function.

III. CONV2D-BASED PHYSICAL NEURAL NETWORK

In this section, we design a physical neural network architecture based on classic Conv2d and propose two realization paradigms: RIS-aided MISO and RIS-aided MIMO systems. Let $\mathbf{X} \in \mathbb{C}^{B \times C_{\text{in}} \times N_w \times N_h}$ and $\mathbf{Y} \in \mathbb{C}^{B \times C_{\text{out}} \times N_w \times N_h}$ denote the input and output matrices of the CNN, respectively, as illustrated in Fig. 1, where B is the batch size, C_{in} and C_{out} denote the numbers of input and output channels, and $N_w \times N_h$ denotes the input dimensions. The convolutional kernels are assumed to have dimensions of $N_k \times N_k$.

A. Conv2d MISO

For MISO systems, i.e., $N_r = 1$, we employ time-division multiple access, where one output channel is received per time slot. Specifically, C_{in} orthogonal frequency division multiplexing (OFDM) carriers are employed at each time, enabling the simultaneous transmission of C_{in} input channels from the transmitter to the receiver. At each time slot t , there are C_{in} kernels $\{\bar{\mathbf{w}}_{i,t}\}$ to emulate, where

$$\bar{\mathbf{w}}_{i,t} = f_{2,i,t} \mathbf{h}_t^H \mathbf{F}_{1,i,t}, \quad i = 1, \dots, C_{\text{in}}, t = 1, \dots, C_{\text{out}}, \quad (2)$$

where $\mathbf{F}_{1,i,t} \in \mathbb{C}^{N_k \times N_k}$ and $f_{2,i,t} \in \mathbb{C}$ denote the i -th OFDM carrier-based precoder at the transmitter and the amplification coefficient at the receiver, respectively. In addition, $\mathbf{h}_t^H = \sum_{l=1}^L \hat{\mathbf{h}}_l^H \Theta_{l,t} \bar{\mathbf{H}}_l$, where $\Theta_{l,i}$ represents the l th RIS phase shift matrix at time slot t .

At the receiver, the outputs of C_{in} channels are piece-wise summed to generate one output channel at each time slot t

$$\mathbf{y}_t = \sum_{i=1}^{C_{\text{in}}} (\bar{\mathbf{w}}_{i,t} \bar{\mathbf{X}} + f_{2,i,t} \mathbf{n}_{i,t}), \quad t = 1, \dots, C_{\text{out}}, \quad (3)$$

where $\mathbf{n}_{i,t}$ denotes the corresponding noise term. After C_{out} transmission time slots, we obtain C_{out} output channels as

TABLE I: Conv2d MISO vs MIMO.

Conv2d	N_t	N_r	T_s	T_r	T_o	T_p	T_c
MISO	N_k^2	1	C_{out}	C_{out}	C_{in}	$C_{out}C_{in}$	0
MIMO	N_k^2	C_{out}	1	1	C_{in}	C_{in}	C_{in}

described in (3). To further enhance transmission efficiency, we dynamically adjust the RIS phase-shift matrix for each time slot, thereby providing more DoFs to emulate the convolution kernels by altering the wireless channels. It is important to note that in this system design, each OFDM carrier is associated with a dedicated precoder at each time slot. Consequently, C_{in} different precoders are employed per time slot, while no combiner is needed at the receiver.

B. Conv2d MIMO

For MIMO systems, we adopt C_{out} receive antennas at the receiver, allowing each antenna to directly capture a distinct output channel. Meanwhile, C_{in} OFDM carriers are still used for transmitting C_{in} input channels to each receive antenna. Specifically, the relationship can be formulated as

$$\bar{\mathbf{W}}_i = \mathbf{F}_{2,i} \mathbf{H} \mathbf{F}_{1,i}, i = 1, \dots, C_{in}, \quad (4)$$

where $\mathbf{F}_{1,i} \in \mathbb{C}^{N_k^2 \times N_k^2}$ and $\mathbf{F}_{2,i} \in \mathbb{C}^{C_{out} \times C_{out}}$ denote the i -th OFDM carrier-based precoder and combiner for convolving with input image $\bar{\mathbf{X}}$, respectively, and $\mathbf{H} = \sum_{i=1}^L \hat{\mathbf{H}}_i \Theta_i \bar{\mathbf{H}}_i$. Since only a single time slot is required for transmission in this setup, the RIS phase-shift matrices need to be adjusted only once, thereby significantly reducing signaling overhead compared to the MISO scheme.

At each receive antenna, C_{in} received channels are summed to generate one output channel as

$$\mathbf{Y} = \sum_{i=1}^{C_{in}} (\bar{\mathbf{W}}_i \bar{\mathbf{X}} + \mathbf{F}_{2,i} \mathbf{N}_i), \quad (5)$$

where $\mathbf{N}_i$ denotes the noise vector associated with the i -th OFDM carrier. Unlike the Conv2d MISO scheme, the Conv2d MIMO design requires the use of C_{in} combiners at the receiver.

The differences between Conv2d MISO and Conv2d MIMO are summarized in Table I. In the table, T_s denotes the number of transmission slots, T_r the number of RIS adjustments, T_o the number of OFDM carriers, T_p the number of precoder adjustments, and T_c the number of combiner adjustments.

IV. CONVSD-BASED PHYSICAL NEURAL NETWORK

In this section, we study the lightweight ConvSD architecture. ConvSD decomposes the convolution process into two stages: depthwise convolution and pointwise convolution [14]. In the depthwise convolution, a single convolution filter is applied per input channel, isolating spatial filtering from inter-channel interactions. In the pointwise convolution, a 1×1 convolution is applied to linearly combine the outputs of all depthwise convolutions across channels, thereby creating new feature representations. As a result, the total number of parameters required for ConvSD is $C_{in} \times N_k^2 + C_{in} \times C_{out}$,

TABLE II: ConvSD MISO vs MIMO.

ConvSD	N_t	N_r	T_s	T_r	T_o	T_p	T_c
MISO	N_k^2	1	1	1	C_{in}	C_{in}	0
MIMO	N_k^2	C_{in}	1	1	1	1	1

which is significantly fewer than that required for a standard Conv2d operation with $C_{in} \times N_k^2 \times C_{out}$ parameters.

A. ConvSD MISO

For the MISO system, C_{in} OFDM carriers are adopted. Each OFDM carrier is assigned a dedicated precoder (i.e., C_{in} different precoders) with each carrier responsible for transmitting one input channel. Since C_{in} OFDM carriers are transmitted simultaneously, only a single transmission slot is needed, and the RIS phase-shift matrix requires adjustment only once. Mathematically, the operation can be expressed as

$$\bar{\mathbf{w}}_i = f_{2,i} \mathbf{h}^H \mathbf{F}_{1,i}, i = 1, \dots, C_{in}, \quad (6)$$

which is structurally similar to (2). After receiving the C_{in} distorted input channels corrupted by noise and channel fading, the receiver applies C_{out} pointwise convolutional filters, each consisting of C_{in} kernels of size 1×1 , in order to emulate the pointwise convolution step. Since $N_r = 1$ in this MISO setting, no combiners are needed at the receiver. The overall operation at the receiver can be formulated as

$$\mathbf{y}_t = \sum_{i=1}^{C_{in}} q_{t,i} (\bar{\mathbf{w}}_i \bar{\mathbf{X}} + f_{2,i} \mathbf{n}_i), t = 1, \dots, C_{out}, \quad (7)$$

where $q_{t,i}$ denotes the i -th kernel coefficient corresponding to the t -th output channel.

B. ConvSD MIMO

For MIMO systems, we set $N_r = C_{in}$, indicating each receive antenna is responsible for one input channel. A combiner of size $C_{in} \times C_{in}$ is adopted at the receiver, jointly designed with the RIS and the precoder to emulate the depthwise convolution. Thus, the overall operation can be expressed as

$$\bar{\mathbf{W}} = \mathbf{F}_2 \mathbf{H} \mathbf{F}_1, \quad (8)$$

which is similar to (4), except that only a single precoder and a single combiner are required in this case. Then, the receiver applies C_{out} filters, each with C_{in} kernels of size 1×1 , to imitate the pointwise convolution process. We have

$$\mathbf{Y}_t = \sum_{i=1}^{C_{in}} q_{t,i} [\bar{\mathbf{W}} \bar{\mathbf{X}} + \mathbf{F}_2 \mathbf{N}]_{i,:}, t = 1, \dots, C_{out}. \quad (9)$$

Note that only one precoder and one combiner are adopted in this case.

A comparison between the parameters of ConvSD MISO and ConvSD MIMO architectures is presented in Table II.




Fig. 3: Multi-RIS aided OTA transmission neural network architecture.




Fig. 4: Transmit power $P_{\max}$ versus classification accuracy.

V. NUMERICAL RESULTS

In this section, we present numerical results to evaluate the image classification accuracy achieved by the proposed schemes, based on the Fashion MNIST dataset. The Fashion MNIST consists of 70,000 images with a resolution of 28×28 pixels, categorized into 10 classes. Among them, 60,000 images are used for training, and 10,000 are utilized for testing. The overall multi-RIS-aided OTA transmission neural network architecture is depicted in Fig. 3, which consists of one convolutional (Conv) layer, one real-to-complex (R2C) layer, two batch normalization (BN) layers, two complex ReLU activation layers, two max pool layers, two complex fully connected (FC) layers, one dropout layer, one complex-to-real (C2R) layer, one real FC layer, and a transceiver module. In addition, the Rician fading channel with Rician factor K is considered for both MIMO and MISO systems. Unless otherwise specified, we set $C_{\text{in}} = 32$, $C_{\text{out}} = 64$, $N_{\text{wi}} = N_{\text{hi}} = 14$, $N_{\text{k}} = 3$, $K = 3$ dB, $N_t = 9$, $N_r = 32$ for ConvSD MIMO, $N_r = 64$ for Conv2d MIMO, $P_{\max} = 10$ dB, and $\sigma^2 = 1$.

In Fig. 4, we study transmit power $P_{\max}$ versus classification accuracy for $L = 1$, $K = 3$ dB, and $M = 100$. The “Upper bound” scheme denotes the digital-domain Conv2d without OTA computation. It can be observed that the classification accuracy of all the schemes except the “Upper bound” increases with $P_{\max}$. This is expected, as higher $P_{\max}$ reduces the detrimental impact of noise at the receiver, thereby gradually approaching the “Upper bound” performance. Additionally, it is observed that the Conv2d-based schemes consistently outperform the ConvSD-based schemes. This is because the




Fig. 5: Rician factor K versus classification accuracy.

ConvSD layer is a simplified version of the Conv2d layer with weaker feature extraction capabilities. Furthermore, Conv2d MISO consistently outperforms Conv2d MIMO. This is because the numbers of adjustments for the precoder and RIS are $C_{\text{in}}C_{\text{out}}$ and C_{out} , respectively, whereas they are C_{in} and 1 for the Conv2d MIMO scheme, meaning that the DoFs available for emulation in the former scheme are much larger. It should be noted that this observation does not hold for the ConvSD scheme. The ConvSD MISO scheme outperforms the ConvSD MIMO scheme only when the transmit power is low, i.e., below 5 dB, but performs worse when the transmit power exceeds 5 dB. This is because the ConvSD MISO scheme adjusts the precoder C_{in} times without adjusting the combiner, whereas the ConvSD MIMO scheme adjusts both the precoder and the combiner once, thus striking different balances.

In Fig. 5, we study Rician factor K versus classification accuracy for $L = 1$, $P_{\max} = 10$ dB, and $M = 50$. As K increases, the channel gain improves, resulting in less signal distortion and enhanced classification accuracy. Furthermore, when K is below -10 dB, the ConvSD MISO scheme outperforms the ConvSD MIMO scheme, however, as K increases, ConvSD MIMO scheme eventually surpasses the ConvSD MISO scheme. This behavior is consistent with the explanation provided for Fig. 4. It is also noteworthy that further increase in K does not always lead to continuous improvements in classification accuracy. Beyond a certain threshold, increasing K may even degrade the performance, as will be discussed later in Fig. 7.

In Fig. 6, we investigate the classification accuracy versus the number of RIS reflecting elements M for $L = 1$,




Fig. 6: Number of reflecting elements M versus classification accuracy.




Fig. 7: Rician factor K versus classification accuracy.

$P_{\max} = 10$ dB, and $K = -10$ dB. It is observed that the classification accuracy of both Conv2D and ConvSD schemes significantly improve with increasing M . This trend can be attributed to two main reasons. First, a larger number of reflecting elements provides more DoFs for modifying the physical neural network. Second, more reflecting elements enhance the beamforming gain, thus mitigating the impact of noise and improving the classification accuracy.

In Fig. 7, we study the classification accuracy of the ConvSD MIMO scheme versus K for $L = 1$, $L = 2$, and $L = 5$ under $P_{\max} = 10$ dB and $M = 50$. It can be observed that for $L = 1$, the classification accuracy initially increases but eventually decreases as K continues to grow. This can be explained as follows. When K is below 15 dB, the channel is dominated by non-line-of-sight (NLoS), resulting in a high channel rank and abundant DoFs for neural network modification. Thus, increasing K improves classification accuracy by enhancing the channel gain. However, when $K \geq 15$ dB, the channel becomes dominated by LoS components. Although the channel gain remains high, the channel rank tends to decrease toward one, limiting the available DoFs, and thereby, degrading the performance. Moreover, we observe that for larger values of L , the system performance remains robust even as K increases. This is because a larger L yields a higher effective channel rank, enhancing the DoFs available for end-to-end training.

VI. CONCLUSION

In this paper, we studied RIS-aided MISO and MIMO systems for engineering the ambient wireless channel to implement CNNs via OTA computation. By jointly training the precoder, combiner, and RIS phase-shift matrices, the digital convolutional operation can be effectively emulated using physical neural networks. We investigated two types of CNNs, namely Conv2d and ConvSD, and proposed two transceiver architectures: RIS-aided MISO and RIS-aided MIMO. A comprehensive comparison between these two architectures for physical CNN implementation was conducted, highlighting the trade-offs between performance gains and implementation overhead. Simulation results demonstrated that the proposed architectures can achieve satisfactory classification accuracy while enabling OTA computation. Furthermore, it was shown that multi-RIS deployments significantly outperform single-RIS when LoS propagation dominates the wireless channel.

REFERENCES

- [1] Q. Wu, B. Zheng, C. You, L. Zhu, K. Shen, X. Shao, W. Mei, B. Di, H. Zhang, E. Basar, L. Song, M. Di Renzo, Z.-Q. Luo, and R. Zhang, "Intelligent surfaces empowered wireless network: Recent advances and the road to 6g," *Proc. IEEE*, vol. 112, no. 7, pp. 724–763, Jul. 2024.
- [2] W. Mei and R. Zhang, "Joint base station and IRS deployment for enhancing network coverage: A graph-based modeling and optimization approach," *IEEE Trans. Wireless Commun.*, vol. 22, no. 11, pp. 8200–8213, Nov. 2023.
- [3] G. Chen, Q. Wu, R. Liu, J. Wu, and C. Fang, "IRS aided MEC systems with binary offloading: A unified framework for dynamic IRS beamforming," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 2, pp. 349–365, Feb. 2023.
- [4] M. Hua, Q. Wu, D. W. K. Ng, J. Zhao, and L. Yang, "Intelligent reflecting surface-aided joint processing coordinated multipoint transmission," *IEEE Trans. Commun.*, vol. 69, no. 3, pp. 1650–1665, Mar. 2021.
- [5] Z. Wang, J. Qiu, Y. Zhou, Y. Shi, L. Fu, W. Chen, and K. B. Letaief, "Federated learning via intelligent reflecting surface," *IEEE Trans. Wireless Commun.*, vol. 21, no. 2, pp. 808–822, Feb. 2022.
- [6] T. Jiang and Y. Shi, "Over-the-air computation via intelligent reflecting surfaces," in *IEEE GLOBECOM, Waikoloa, HI, USA*, 2019, pp. 1–6.
- [7] E. Arslan, I. Yildirim, F. Kilinc, and E. Basar, "Over-the-air equalization with reconfigurable intelligent surfaces," *IET Commun.*, vol. 16, no. 13, pp. 1486–1497, May 2022.
- [8] C. Liu, Q. Ma, Z. J. Luo, Q. R. Hong, Q. Xiao, H. C. Zhang, L. Miao, W. M. Yu, Q. Cheng, L. Li et al., "A programmable diffractive deep neural network based on a digital-coding metasurface array," *Nat. Electron.*, vol. 5, no. 2, pp. 113–122, Feb. 2022.
- [9] S. Chen, Y. Hui, Y. Qin, Y. Yuan, W. Meng, X. Luo, and H.-H. Chen, "RIS-based on-the-air semantic communications – a diffractive deep neural network approach," *IEEE Wireless Commun.*, vol. 31, no. 4, pp. 115–122, Aug. 2024.
- [10] K. Stylianopoulos, P. Di Lorenzo, and G. C. Alexandropoulos, "Over-the-air edge inference via end-to-end metasurfaces-integrated artificial neural networks," 2025. [Online]. Available: <https://arxiv.org/abs/2504.00233>.
- [11] Y. Yang, Z. Zhang, Y. Tian, Z. Yang, R. Jin, L. Liu, and C. Huang, "Realizing over-the-air neural networks in RIS-assisted MIMO communication systems," in *IEEE WCNC, Dubai, United Arab Emirates*, 2024, pp. 1–5.
- [12] S. Garcia Sanchez, G. Reus-Muns, C. Bocanegra, Y. Li, U. Muncuk, Y. Naderi, Y. Wang, S. Ioannidis, and K. R. Chowdhury, "AirNN: Over-the-air computation for neural networks via reconfigurable intelligent surfaces," *IEEE/ACM Tran. Netw.*, vol. 31, no. 6, pp. 2470–2482, Dec. 2023.
- [13] J. Zhang, H. Chen, and D. M. Blough, "A radio-frequency-based 2-D convolutional layer using transmissive intelligent surfaces," in *IEEE VTC, Washington, DC, USA*, 2024, pp. 1–7.
- [14] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1251–1258.