

# Zero-Shot Adaptation for mmWave Beam-Tracking on Overhead Messenger Wires through Robust Adversarial Reinforcement Learning

Masao Shinzaki, *Student Member, IEEE*, Yusuke Koda, *Graduate Student Member, IEEE*, Koji Yamamoto, *Senior Member, IEEE*, Takayuki Nishio, *Senior Member, IEEE*, Masahiro Morikura, *Member, IEEE*, Yushi Shirato, *Member, IEEE*, Daisei Uchida, *Member, IEEE*, and Naoki Kita, *Member, IEEE*

**Abstract**—Millimeter wave (mmWave) beam-tracking based on machine learning enables the development of accurate tracking policies while obviating the need to periodically solve beam-optimization problems. However, its applicability is still arguable when training-test gaps exist in terms of environmental parameters that affect the node dynamics. From this skeptical point of view, the contribution of this study is twofold. First, by considering an example scenario, we confirm that the training-test gap adversely affects the beam-tracking performance. More specifically, we consider nodes placed on overhead messenger wires, where the node dynamics are affected by several environmental parameters, e.g., the wire mass and tension. Although these are particular scenarios, they yield insight into the validation of the training-test gap problems. Second, we demonstrate the feasibility of *zero-shot adaptation* as a solution, where a learning agent adapts to environmental parameters unseen during training. This is achieved by leveraging a robust adversarial reinforcement learning (RARL) technique, where such training-and-test gaps are regarded as disturbances by adversaries that are jointly trained with a legitimate beam-tracking agent. Numerical evaluations demonstrate that the beam-tracking policy learned via RARL can be applied to a wide range of environmental parameters without severely degrading the received power.

**Index Terms**—mmWave communications, beam-tracking, robust adversarial reinforcement learning, zero-shot adaptation, overhead messenger wire.

## I. INTRODUCTION

WIRELESS communication technologies in fifth-generation (5G) mobile networks provide multigigabit-per-second data rates, which fulfill the backhaul rate requirements [1], [2]. A key technology of 5G systems is millimeter-wave (mmWave) communications, which is advantageous because its broad spectral band increases the communication capacity [3]. In contrast to time- and cost-intensive optical fiber deployments, a mmWave wireless backhaul network has the advantages of high flexibility, cost efficiency, and rapid deployment of backhaul connections [1].

Masao Shinzaki, Koji Yamamoto, and Masahiro Morikura are with the Graduate School of Informatics, Kyoto University, Kyoto 6068501 Japan. e-mail: {shinzaki@imc.cce., kyamamoto@i.kyoto-u.ac.jp, and morikura.masahiro.h21@kyoto-u.jp.

Yusuke Koda is with the Centre for Wireless Communications, University of Oulu, 90014, Finland, e-mail: Yusuke.Koda@oulu.fi.

Takayuki Nishio is with School of Engineering, Tokyo Institute of Technology, Ookayama, Meguro-ku, Tokyo, 158-0084, Japan, e-mail: nishio@ict.e.titech.ac.jp.

D. Uchida, Y. Shirato, and N. Kita are with the NTT Access Network Service Systems Laboratories, NTT Corporation, Hikarinooka, Yokosuka 2390847 Japan.

However, signals transmitted in the mmWave frequency band experience larger path loss, which mandates the use of large antenna arrays and adaptive control of the array weights to ensure the main lobes of the antennas are pointing toward each other [4]. This adaptive control of the antenna array is required not only for initial access, but also during an operation in a mobile scenario or quasi-static scenarios where mmWave nodes are occasionally displaced owing to various perturbations. The latter is termed beam-tracking, and in view of the large overhead involved in tracking, developing efficient beam-tracking methods has attracted considerable research interest, as is briefly discussed below.

### A. Related Work and Motivations

Previous research addressed the beam-tracking problem mainly via the following two approaches: 1) periodically solving beam-optimization problems [5]–[11], 2) learning a beam-tracking policy beforehand [12]–[21]. In the first approach, a mmWave node optimizes the array antenna weights based on the estimated channels or angles of arrival (AoAs)/angles of departure (AoDs). For example, in the IEEE 802.11ad standard [5], channels are surveyed by steering the transmitter/receiver beams; thus the array antenna weights are optimized. The work in [6] estimated the AoAs/AoDs, and optimized the array antenna weights based on these estimations. The works in [7]–[11] used filtering methods to estimate the AoAs/AoDs based on the surveyed channel and calculated the optimal antenna weights. Although adaptive, this approach incurs computational overhead to periodically solve optimization problems that scale a number of antennas.

In the second approach, with the help of powerful machine-learning (ML) techniques approximating the input-output relationships, appropriate antenna weights or beam steering angles are learned beforehand. Although the computational overhead increases during the training procedure, appropriate antenna weights or beam-steering angles can be outputted with fewer computations than solving the optimization problems thereafter. Prior studies mainly leveraged supervised learning (SL), where the training data for optimal angles are given in advance [12]–[17]. Wang *et al.* and Mismar *et al.* [18], [19] leveraged reinforcement learning (RL) techniques, where a mmWave node learns appropriate beam steering angles from the received power without being given the optimal angles as training data. Our previous studies relating to this work

TABLE I  
COMPARISON FROM PREVIOUS WORKS ON BEAM-TRACKING

| Reference                   | 1) Periodically solving beam-optimization? | 2) Learning-Based? | Addressing Training-Test Gap? |
|-----------------------------|--|--------------------|-------------------------------|
| [5]–[11]                    | Yes  | No                 | –                             |
| [12]–[17]                   | No   | Yes (SL)           | No                            |
| [18], [19]                  | No   | Yes (RL)           | No                            |
| Previous version [20], [21] | No   | Yes (RL)           | No                            |
| This paper                  | No   | Yes (RL)           | Yes                           |

[20], [21] were also categorized as focusing on RL-based beam-tracking, where we considered the beam-tracking of a mmWave node placed on an overhead messenger wire.

Although this second approach is attractive with respect to its ability to perform beam-tracking, its applicability to the real environment is arguable because this approach may experience a training-test gap in terms of the environmental parameters that affect the node dynamics, which is the main focus of this study. Generally, ML models perform worse as the gap between training and testing increases in terms of dataset distributions in SL or parameters that determine the state dynamics in RL. Hence, this performance deterioration may naturally occur in the above learning-based beam-tracking when a training-test gap exists. To the best of our knowledge, this problem has been overlooked in the aforementioned prior studies pertaining to learning-based beam-tracking. Note that this is also the main difference from our prior studies [20], [21], where the training-test gap affecting node dynamics was not considered; in this sense, the contribution is different from these previous studies. The differences between this work and the previous works are summarized in Table I.

### B. Contributions

In view of the above problem, this work addresses the following two questions as the contribution of this study: 1) what would happen if a training-test gap were to exist in learning-based beam-tracking? 2) If this gap has a negative effect on beam-tracking, how could we solve the problem caused by this training-test gap? To address the first question, we confirm that the training-test gap causes the received power to deteriorate via numerical evaluation. To address the second question, we first conceptualize *zero-shot adaptation* as our objective. Here, zero-shot adaptation implies that a learning agent exhibits feasible performance in test scenarios even when there are training-test gaps. To realize zero-shot adaptation in mmWave beam-tracking, we applied robust adversarial reinforcement learning (RARL) [22], which is detailed below. It should be noted that when examining these problems, we consider a particular scenario in which a mmWave node is placed on an overhead messenger wire, similar to our previous work ([20] and [21]). This is because this scenario is affected by several parameters that determine the node dynamics, such as the mass and tension of the wire, which facilitate the investigation of the training-test gap problem. We believe that, although we are considering a particular scenario, this work provides firsthand insight regarding the overlooked problem caused by the training-test gap and provides the opportunity to

rethink learning-based systems in the wireless communication research area. In view of this, the contributions of this work are summarized as follows:

- Given an example scenario in which a mmWave node is placed on an overhead messenger wire (see Fig. 2), through numerical evaluations, we confirm that the gap between the training and test scenarios deteriorates the beam-tracking performance. This debunks the importance of addressing training-test gaps to provide reliable mmWave links, and to the best of our knowledge, this perspective has not yet been reported in the literature.
- We demonstrate the feasibility of zero-shot adaptation of learning-based mmWave beam-tracking in the aforementioned scenarios. The key idea is to leverage RARL, wherein a beam-tracking agent is trained competitively to correspond to an intelligent adversary that attempts to cause beam misalignment by introducing additional wind disturbance. Through numerical evaluations, we show that even if the test scenarios were different from the training scenarios in terms of the parameters that affect the node dynamics, the proposed method prevents a drastic performance loss in terms of received power without adaptively fine-tuning the test scenario.

In a nutshell, our main scope is to validate the effectiveness of adding an intelligent adversary during the training; thereby confirming the concept of the zero-shot adaptation in the context of mmWave beam-tracking. Thus, particularly in the evaluation in Section V, we focus on the difference between the proposed method with the adversary and a baseline method without the adversary. For this reason, delving into the problems commonly applied to the RL-based beam-tracking method with and without the adversary (e.g., delay for tracking and the accuracy of state acquisitions) is beyond the scope of this paper. We believe that without addressing these problems, our evaluations sufficiently validate the aforementioned contributions.

It should be also noted that the training-test gap in terms of the environmental parameters affecting the node dynamics may be an issue not only in beam-tracking problems but also in other problems in mmWave networks, e.g., a handover problem in mobile scenarios. However, we focus only on beam-tracking problems in a non-mobile scenario because we believe that this is sufficient to answer the above two questions and make the two contributions summarized above. Hence, providing a concrete design on how to overcome the training-test gap issue in the other scenarios and problems is deferred as a future challenge, which is excluded in the technical part

of this paper and is briefly discussed in Section VI.

The remainder of this paper is organized as follows. In Section II, we introduce zero-shot adaptation to overcome the training-test gap in RL problems and provide the motivation for solving this training-test gap in the mmWave beam-tracking problem. In Section III, we formulate the system model as an RL task. In Section IV, we explain the adversarial RL algorithm. In Section V, we describe the simulation evaluation of the proposed beam-tracking policy. Finally, we present our conclusions in Section VI.

## II. DEFINITION AND MOTIVATION OF ZERO-SHOT ADAPTATION

### A. Definition of Zero-Shot Adaptation

We define the zero-shot adaptation problem considered in this study. Let us consider a Markov decision process (MDP)  $(\mathcal{S}, \mathcal{A}, r, p_{\theta})$ , where  $\mathcal{S}$  and  $\mathcal{A}$  denote the state and action space, respectively,  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is the reward function, and  $p_{\theta} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  is the state transition rule. Note that the state transition rule is subject to the static parameter  $\theta$ , which we refer to as the environmental parameters. Therein, at each time step  $k = 1, 2, \dots, K$ , a decision maker observes a state  $s_k \in \mathcal{S}$  and determines an action  $a_k \in \mathcal{A}$  according to a policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$ , and receives a reward  $r_k$ . The objective of the decision-maker is to seek for a policy that maximizes the expected sum of rewards  $\sum_{k=1}^K \gamma^{k-1} r_k$ , where  $\gamma \in [0, 1]$  is the discount factor. Given the aforementioned decision process, we define the zero-shot adaptation problem as follows:

**Definition 1** (Zero-shot adaptation). *Let us consider two MDPs  $MDP_{\text{train}} = (\mathcal{S}, \mathcal{A}, r, p_{\theta_{\text{train}}})$  and  $MDP_{\text{test}} = (\mathcal{S}, \mathcal{A}, r, p_{\theta_{\text{test}}})$  are different in terms of the environmental parameters, that is,  $\theta_{\text{train}} \neq \theta_{\text{test}}$ . During the learning procedure of a policy, the decision-maker can act only in  $MDP_{\text{train}}$  and cannot access  $\theta_{\text{test}}$ . At the same time, the decision-maker or another third-party agent can “manipulate” training data, that is, they can replace several elements in the obtained state-action trajectory  $(s_1, a_1, \dots, s_K, a_K)$  with other values. Given this constraint, zero-shot adaptation is defined as finding a policy that maximizes the expected sum of rewards in  $MDP_{\text{test}}$  without any re-training of the policy.*

In a nutshell, the zero-shot adaptation considered in this study is to find the optimal policy in an unseen environment for a decision-maker with the tolerance of manipulating the state-action trajectory. Hence, the problem reduces the manipulation of the state-action trajectory, that is, the training data, where the RARL is an effective approach, as demonstrated throughout this study.

Note that this problem is not necessarily identical to a “zero-shot learning problem” [23] in an SL context in terms of the way in which the problem regarding the training-test gap is overcome. In the zero-shot learning as defined in [23], an SL model is trained to enable the model to classify not only samples with a class label seen during training, but also those with unseen classes during training. Rather than achieving this by manipulating the training data, it is accomplished with “auxiliary knowledge,” which indicates

pre-obtained feature information involving class labels unseen in training<sup>1</sup>. However, we also consider our problem to be “zero-shot” by focusing on the common aspect underlying both problems concerned with the training-test gap.

### B. Motivation of Zero-Shot Adaptation in mmWave Beam-Tracking

As is shown in the subsequent sections, we consider the beam-tracking problem on an overhead messenger wire. This scenario involves several environmental parameters, and among these, we select an overall mass  $m$  and spring constant (i.e., wire tension)  $k_0$  of the wire to which a mmWave node is attached. These parameters are hardly measured precisely, particularly when a messenger wire is pre-installed. Hence, to attach the mmWave nodes to such a pre-installed messenger wire, a beam-tracking policy should be trained without accessing these parameters.

One possible approach is to train the beam-tracking policy after attaching the nodes to a pre-installed messenger wire such that  $\theta_{\text{train}} \simeq \theta_{\text{test}}$ . However, this approach has the following two drawbacks, both of which can be solved via zero-shot adaptation. First, in this approach, because it is necessary to train the beam-tracking policy after the attachment, supplying the connections is delayed. However, if the beam-tracking policy could be pre-trained by simulations via zero-shot adaptation, the connection could be supplied immediately, which would be preferable for real deployments. Second, even if the beam-tracking policy could be trained quickly using the first approach, the aforementioned parameters gradually but certainly vary with time owing to the degradation of the wire over time. This yields another test scenario  $\theta'_{\text{test}}$ , where  $\theta_{\text{train}} \neq \theta'_{\text{test}}$ . Hence, to obtain robustness against the yielded training-test gap, it is worthwhile to consider zero-shot adaptation.

However, it remains unclear whether the training-test gap is harmful in the context of mmWave beam-tracking. Hence, for the sake of clarity, we conclude this section by displaying a partial result in Section V-D, as shown in Fig. 1, which shows what happens if there exists a training-test gap. The beam-tracking policy was trained with the parameter setting indicated by a cross mark, that is, a wire mass of 10 kg and wire tension 100 N/m; namely,  $\theta_{\text{train}} = [10 \text{ kg}, 100 \text{ N/m}]^T$ . We introduce this beam-tracking policy in other parameter settings  $\theta_{\text{test}} = [m_{\text{test}}, k_{\text{test}}]^T$ , which are depicted as the horizontal and vertical axes in Fig. 1, respectively. Hence, apart from the cross mark point, there exists a training-test gap, that is,  $\theta_{\text{test}} \neq \theta_{\text{train}}$ . As shown in Fig. 1, under several settings of  $\theta_{\text{test}}$ , the learned beam-tracking policy does not perform better than under  $\theta_{\text{train}}$  in terms of the received power. In particular, the learned beam-tracking policy exhibited poorer received power when the wire mass was lower than that in the test scenario. This is because a wire with a smaller mass has more vibrant wire dynamics, which makes beam-tracking more challenging than

<sup>1</sup>A well-known example is the problem of classifying an image of a “zebra.” Even if an image of the zebra is not contained in the training data, it would be possible to predict that it is an image of a zebra when it is known that a zebra has the appearance of a “striped horse,” and images labeled with “horse” and “striped” appear in the training dataset.

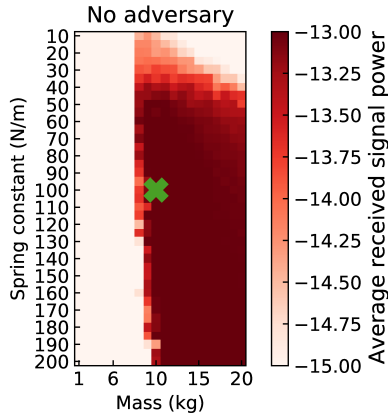


Fig. 1. Heatmap depicting the robustness against training-test gap without zero-shot adaptation in terms of the received power. The cross mark point  $\theta_{\text{train}} = [10\text{kg}, 100\text{N/m}]^T$  represents the environmental parameters of the total wire mass and spring constant (i.e., wire tension) during training. The learned policy was brought in the different environmental parameter settings  $\theta_{\text{test}}$ , where the wire mass and spring constant are depicted in horizontal and vertical axes, respectively.

that in the training scenario. This example led us to address the aforementioned zero-shot adaptation problem in the context of mmWave beam-tracking on a messenger wire.

### III. SYSTEM MODEL

Fig. 2 shows the system model for an on-wire small-cell base station (SBS) mmWave backhaul connection. The SBS on the overhead messenger wire communicates with the gateway BS mounted on the building surface through a mmWave link to relay data from the gateway BS to the overhead messenger wire, and vice-versa. The SBS is also connected to the optical fibers installed along the wire physically and receives the data to be transferred to the gateway BS via the upper layer protocol rather than the physical (PHY) layer (e.g., internet protocol). A possible configuration for data delivery is as follows. Both the SBS and gateway BS function as routers that forward the data from/to the end user, who sits in the house in Fig. 2. The SBS hereby receives data routed to the end user via the upper layer protocol because the SBS is the natural waypoint for data delivery to the end user. Briefly, the SBS provides a connection to the Internet infrastructure directly from an overhead messenger wire near the building. However, it should be noted that the problem of beam-tracking that we consider in this study lies in the PHY layer, which is different from the data delivery problem in terms of the OSI reference model. Hence, we focus on the former problem without a specific consideration for the upper layer protocols, which is sufficient to validate the effectiveness of the proposed beam-tracking method. In this system, the SBS installed on an overhead messenger wire with a weight of  $m$  performs beamforming to increase the received signal power of the gateway BS. The endpoints of the overhead messenger wire are fixed to the telephone poles at a height  $h_w$ , with a distance  $d_w$  between the poles. The gateway BS is mounted at a height  $h_r$ , and the distance between the overhead messenger wire and the gateway is  $d_r$ .

TABLE II  
SUMMARY OF NOTATIONS IN SYSTEM MODEL

|  |   |
|--|---|
| $h_w$                                    | Height of endpoints of wire                               |
| $d_w$                                    | Direct distance between endpoints                         |
| $h_r$                                    | Height of gateway BS                                      |
| $d_r$                                    | Distance between wire and gateway BS                      |
| $m$                                      | Total wire mass   |
| $N$                                      | Number of proxy points on wire                            |
| $i$                                      | Index of proxy point on wire                              |
| $k_0$                                    | Spring constant, i.e., coefficient for tensile force [24] |
| $g$                                      | Gravitational acceleration                                |
| $\mathbf{a}_i(t)$                        | Acceleration of proxy point $i$ at time $t$               |
| $\mathbf{v}_i(t)$                        | Velocity of proxy point $i$ at time $t$                   |
| $\mathbf{x}_i(t)$                        | Position of proxy point $i$ at time $t$                   |
| $c_0$                                    | Drag constant [25]  |
| $\mathbf{v}_0(t)$                        | Wind velocity   |
| $\mathbf{V}_0$                           | Covariance matrix of wind velocity                        |
| $\mathbf{W}_i(t)$                        | Standard Wiener process                                   |
| $P_t$                                    | Transmit power  |
| $\lambda$                                | Radio-wave wavelength                                     |
| $\theta_{\text{AoD}}, \phi_{\text{AoD}}$ | Arbitrary zenith/azimuth angles                           |
| $\theta_s, \phi_s$                       | Zenith/azimuth angles of antenna main-lobe                |
| $d$                                      | Distance between SBS and gateway BS                       |
| $A_r$                                    | Receiver antenna gain                                     |
| $G_{\text{max}}$                         | Antenna gain of main-lobe                                 |
| $A_m$                                    | Front-back ratio [26]                                     |
| $\theta_{3\text{dB}}, \phi_{3\text{dB}}$ | Horizontal/vertical 3 dB beamwidth                        |
| $SLA_V$                                  | Side-lobe level limit [26]                                |
| $n_v, n_H$                               | Number of vertical/horizontal array elements              |
| $\mathbf{w}$                             | Beamforming vector  |
| $\Delta_v, \Delta_H$                     | Vertical/horizontal array spacing distances               |
| $\tau$                                   | Decision interval, i.e., interval between time steps      |
| $k$                                      | Index of time step  |
| $T$                                      | Observation time  |
| $\mathbf{x}_S^{(k)}$                     | Position of SBS at time step $k$                          |
| $\mathbf{x}_G$                           | Position of gateway BS                                    |

For practical usage, beam-tracking on the side of both the gateway and the SBS should also be addressed; nonetheless, we consider only SBS-side beam-tracking. This is because of our focus on validating the effectiveness of adding the adversary during training, thereby confirming the feasibility of the zero-shot adaptation. Indeed, gateway-side beam-tracking is more challenging than SBS-side beam-tracking in the sense that the gateway BS cannot immediately obtain the necessary state information in the RL-based beam-tracking (i.e., position/velocity of the SBS). However, this challenge is equally posed to beam-tracking trained both with and without an adversary, indicating that the difference between gateway-side and SBS-side beam-tracking does not have a specific impact on the comparison between these two beam-tracking methods (i.e., with or without adversaries). Hence, to focus on the comparison, we consider only SBS-side beam-tracking by assuming that the gateway-side beam-tracking was performed perfectly, which is sufficient to validate our contributions.

#### A. Model of Dynamics in Overhead Messenger Wire

According to [24], we modeled the overhead messenger wire as a chain of several proxy mass points that align and are separated by an equal distance, where the proxy mass points are affected by a tensile force from adjacent mass points. Let  $N$  and  $m$  denote the number of proxy mass points and the total mass of the wire, respectively. In the model, the mass of each point is assigned equally, that is, the mass of each point is




Fig. 2. On-wire SBS in a millimeter-wave backhaul connection.




Fig. 3. Coordinate system of the system model.

$m/N$ . We denote points 1 and  $N$  as the ends of the chain of the mass points and term the residual mass points in the order of their proximity to point 1 as point 2, ..., point  $N-1$ . In the model, the tensile force applied to point  $i \in \{2, \dots, N-1\}$  is proportional to the relative position of the adjacent points, i.e., points  $i-1$  and  $i+1$ , where the total tensile force is calculated as:  $k_0[x_{i+1}(t) + x_{i-1}(t) - 2x_i(t)]$ , where  $x_i(t)$  is the position of point  $i$  at time  $t$  measured in the coordinate system in Fig. 3. The term  $k_0$  is constant and determines a wire tension. This model can be regarded as a spring chain, where the mass points are connected via springs; hence, we refer to the constant  $k_0$  as the “spring constant,” hereinafter.

We delve into the dynamics of these proxy mass points. Let  $\mathbf{a}_i(t) \in \mathbb{R}^3$  denote the acceleration of point  $i$  at time  $t$ . The accelerations of points 1 and  $N$  fixed to the telephone poles are expressed as  $\mathbf{a}_1(t) = \mathbf{a}_N(t) = \mathbf{0}$ . For  $i = 2, 3, \dots, N-1$ , from the equation of motion,  $\mathbf{a}_i(t)$  is given by

$$\mathbf{a}_i(t) = \mathbf{g} + \underbrace{\frac{k_0 N}{m} [x_{i+1}(t) + x_{i-1}(t) - 2x_i(t)]}_{\text{acceleration from tensile force}}, \quad (1)$$

where  $\mathbf{g} \in \mathbb{R}^3$  denotes the gravitational acceleration.

As the perturbations that are responsible for the dynamics in the wire, we consider a wind perturbation and consider the wind drag model in [25]. In this model, the wind drag primarily consists of frictional drag and pressure drag [25]. The frictional drag increases proportionally with the velocity of the mass point of interest  $\mathbf{v}_i(t)$  relative to the wind velocity

$\mathbf{v}_0(t)$ . We denote the constant of the proportionality as  $c_0$  and refer to it as the “drag constant.” The pressure drag has random magnitude regardless of the wind speed. The derivatives of the velocity and position of point  $i$  at time  $t$  are denoted by  $d\mathbf{v}_i(t) \in \mathbb{R}^3$  and  $d\mathbf{x}_i(t) \in \mathbb{R}^3$ , respectively. The derivatives of the velocities and positions of points 1 and  $N$  are expressed as  $d\mathbf{v}_1(t) = d\mathbf{v}_N(t) = d\mathbf{x}_1(t) = d\mathbf{x}_N(t) = \mathbf{0}$ , respectively. For  $i = 2, 3, \dots, N-1$ ,  $d\mathbf{v}_i(t)$  and  $d\mathbf{x}_i(t)$  are calculated as follows [25], [27]:

$$\begin{aligned} d\mathbf{v}_i(t) &= \mathbf{a}_i(t) dt - \underbrace{c_0 [\mathbf{v}_i(t) - \mathbf{v}_0(t)] dt + \mathbf{V}_0 d\mathbf{W}_i(t)}_{\text{derivatives from the wind}}, \\ d\mathbf{x}_i(t) &= \mathbf{v}_i(t) dt, \end{aligned} \quad (2)$$

where  $\mathbf{V}_0 \in \mathbb{R}^{3 \times 3}$  and  $\mathbf{W}_i(t) \in \mathbb{R}^3$  denote the covariance matrix of the wind speed and the standard Wiener process that is independently and identically distributed across positions of point  $i$ , respectively.

### B. Radio-Wave Propagation Model

According to the free-space path loss model [28], we consider the received signal power of the gateway BS to be determined by the distance  $d$  between the SBS and the gateway BS and the antenna radiation pattern. This free-space path loss model is consistent with the mmWave channel measurement conducted under a line-of-sight (LoS) conditions in an open space [29], and this is a feasible assumption considering that the SBS and gateway BS in the above-mentioned scenario are likely to be deployed under such conditions. Note that in a NLoS condition, the above assumption does not hold indeed; however, we do not consider the NLoS condition because our focus is on the beam-tracking inaccuracy caused by the training-test gap, which occurs even under LoS conditions. In other words, considering only the LoS condition is sufficient to validate the contributions of this study. We consider the use of a directional antenna; hence, the antenna radiation pattern is determined by AoDs in the zenith and azimuth angles. That is,  $\theta_{\text{AoD}}, \phi_{\text{AoD}}$ , respectively, and the zenith and azimuth steering angles, that is,  $\theta_s, \phi_s$ , respectively. Note that these angles are measured in the coordinate system illustrated in Fig. 3. From the Friis transmission equation [28],  $P_r(d, \theta_{\text{AoD}}, \phi_{\text{AoD}}, \theta_s, \phi_s)$  is given by

$$P_r(d, \theta_{\text{AoD}}, \phi_{\text{AoD}}, \theta_s, \phi_s) = \left( \frac{\lambda}{4\pi d} \right)^2 A_t(\theta_{\text{AoD}}, \phi_{\text{AoD}}, \theta_s, \phi_s) A_r P_t, \quad (3)$$

where  $P_t$ ,  $\lambda$ , and  $A_r$  are constants, and denote the transmission power of the SBS, wavelength of the radio waves, and receiver antenna gain, respectively. Moreover,  $A_t(\theta_{\text{AoD}}, \phi_{\text{AoD}}, \theta_s, \phi_s)$  denotes the transmission antenna gain, with its maximum value at  $\theta_{\text{AoD}} = \theta_s$  and  $\phi_{\text{AoD}} = \phi_s$ . For the sake of simplicity, we omit the subscript  $\text{AoD}$ , hereinafter.

We considered the use of the array antenna model in [26], [30], where the transmission antenna gain  $A_t(\theta, \phi, \theta_s, \phi_s)$  is given by

$$A_t(\theta, \phi, \theta_s, \phi_s) = A_E(\theta, \phi) + A_F(\theta, \phi, \theta_s, \phi_s), \quad (4)$$

where  $A_E(\theta, \phi)$  and  $AF(\theta, \phi, \theta_s, \phi_s)$  denote the element radiation pattern and array factor, respectively. The element radiation pattern  $A_E(\theta, \phi)$  of each single antenna element is composed of horizontal and vertical radiation patterns. The element radiation pattern  $A_E(\theta, \phi)$  is given by

$$A_E(\theta, \phi) = G_{\max} - \min \left\{ -[A_{E,V}(\theta) + A_{E,H}(\phi)], A_m \right\}, \quad (5)$$

where  $A_{E,V}(\theta)$ ,  $A_{E,H}(\phi)$ ,  $G_{\max}$ , and  $A_m$  denote the vertical and horizontal radiation patterns, maximum directional gain of the antenna element, and front-back ratio, respectively. The vertical and horizontal radiation patterns  $A_{E,V}(\theta)$  and  $A_{E,H}(\phi)$ , respectively, are obtained as follows:

$$\begin{aligned} A_{E,V}(\theta) &= -\min \left\{ 12 \left( \frac{\theta - 90^\circ}{\theta_{3dB}} \right)^2, SLA_V \right\}, \\ A_{E,H}(\phi) &= -\min \left\{ 12 \left( \frac{\phi}{\phi_{3dB}} \right)^2, A_m \right\}, \end{aligned} \quad (6)$$

where  $\theta_{3dB}$ ,  $\phi_{3dB}$ , and  $SLA_V$  are the vertical 3 dB beamwidth, horizontal 3 dB beamwidth, and side-lobe level limit, respectively.

The array factor  $AF(\theta, \phi, \theta_s, \phi_s)$  models the directivity of the antenna array, which is expressed for an array of  $n = n_V n_H$  elements as

$$AF(\theta, \phi, \theta_s, \phi_s) = 10 \log_{10} \left[ 1 + \left( (1/\sqrt{n})w \right)^2 - 1 \right], \quad (7)$$

where  $n_V$  and  $n_H$  denote the number of vertical and horizontal elements, respectively, and  $w \in \mathbb{C}^n$  denotes the beamforming vector, which is given by

$$\begin{aligned} w &= [w_{1,1}, w_{1,2}, \dots, w_{n_V, n_H}]^T, \\ w_{p,r} &= e^{j2\pi[(p-1)\Delta_V \Psi_p / \lambda + (r-1)\Delta_H \Psi_r] / \lambda}, \\ \Psi_p &= \cos \theta - \cos \theta_s, \\ \Psi_r &= \sin \theta \sin \phi - \sin \theta_s \sin \phi_s, \end{aligned} \quad (8)$$

where  $\Delta_V$  and  $\Delta_H$  denote the spacing distances between the vertical and horizontal elements of the array, respectively. As a specific characteristic of mmWave communications, the spacing distances  $\Delta_V$  and  $\Delta_H$  are of the order of several millimeters (e.g.,  $\Delta_V = \Delta_H = 2.5$  mm at an the RF frequency of 60 GHz). This is because of the common setting where spacing distances should not exceed the half-wavelength to avoid high grating lobes.

The SBS periodically observes both the instantaneous received signal power and its position and velocity. Hereinafter, we let the notation  $\tau$  denote the observation interval and term the time instants for the observation as the “time step.” Accordingly, we consider the steering angle capable of moving up, down, left, or right by an angle  $\beta$  at each time step. The problem for determining the steering angles at each time step is formulated in the next section.

### C. Initial Access Procedure

Among the initial access procedures, we only note the beam alignment between the SBS and gateway BS in the initial stage. Indeed, many other procedures are mandatory to initialize the mmWave communications (e.g., device discovery

and association frame exchanges); however, our focus is on beam-tracking, which is disjoint these initial access procedures. Hence, providing a concrete design of the initial access procedures is basically beyond the scope of this paper, and we describe only the beam alignment in the initial stage, which is the most relevant procedure for beam-tracking.

The beam alignment in the initial stage should allow the beam of the SBS to point toward the gateway BS; thereby allowing the gateway BS to receive the maximum received power. Indeed, one can assume arbitrary beam-alignment procedures as long as the beam of the SBS to point toward the gateway BS. As an example of this procedure, in the evaluation in Section V, we employed position-aware beam alignment based on our assumption that the positions of the SBS and gateway BS are given in advance. Therein, we simply calculated the orientation of the gateway BS from these positions under the initial conditions without wind. Subsequently, we established the array antenna weights such that the beam pointed toward the gateway BS.

### D. Formulation

Let  $\mathcal{N} = \{1, 2, \dots, \lfloor T/\tau \rfloor\}$  denote the set of indices of the time steps, where the index  $k \in \mathcal{N}$  corresponds to the time  $t = k\tau$ . The term  $T$  is the total time length for the beam-tracking. Moreover, we let the superscript  $(k)$  indicate that the variables of interest are measured at the time  $t = k\tau$ . The optimization problem is formulated as follows:

$$\begin{aligned} &\underset{(a_\theta^{(k)}, a_\phi^{(k)})_{k \in \mathcal{N}}}{\text{maximize}} && \frac{1}{\lfloor T/\tau \rfloor} \sum_{k \in \mathcal{N}} P_r(d^{(k)}, \theta^{(k)}, \phi^{(k)}, \theta_s^{(k)}, \phi_s^{(k)}). \end{aligned} \quad (9a)$$

subject to

$$d^{(k)} = \|\mathbf{x}_S^{(k)} - \mathbf{x}_G\|, \quad (9b)$$

$$\theta^{(k)} = \arccos \frac{x_{S,z}^{(k)} - x_{G,z}}{d^{(k)}}, \quad (9c)$$

$$\phi^{(k)} = \arctan \frac{x_{S,y}^{(k)} - x_{G,y}}{x_{S,x}^{(k)} - x_{G,x}}, \quad (9d)$$

$$\theta_s^{(k)} = \theta_s^{(k-1)} + a_\theta^{(k)} \beta, \quad (9e)$$

$$\phi_s^{(k)} = \phi_s^{(k-1)} + a_\phi^{(k)} \beta, \quad (9f)$$

$$a_\theta^{(k)}, a_\phi^{(k)} \in \{-1, 0, 1\}, \quad (9g)$$

$$|a_\theta^{(k)}| + |a_\phi^{(k)}| \leq 1, \quad (9h)$$

where  $a_\theta^{(k)}$  and  $a_\phi^{(k)}$  denote the action for the zenith and azimuth angles at time step  $k$ , respectively, which can move the zenith and azimuth steering angles  $\theta_s^{(k)}, \phi_s^{(k)}$  by angle  $\beta$ , respectively. Moreover,  $d^{(k)}$ ,  $\theta^{(k)}$ , and  $\phi^{(k)}$  denote the distance, beam zenith, and azimuth angle from the SBS to the gateway BS at time step  $k$ , respectively. These variables are defined by the positions of the SBS and gateway BS, i.e.,  $\mathbf{x}_S^{(k)} = [x_{S,x}^{(k)}, x_{S,y}^{(k)}, x_{S,z}^{(k)}]^T \in \mathbb{R}^3$ ,  $\mathbf{x}_G = [x_{G,x}, x_{G,y}, x_{G,z}]^T \in \mathbb{R}^3$ , respectively.



#### IV. ADVERSARIAL RL-BASED BEAM-TRACKING BASED ON ZERO-SHOT ADAPTATION

##### A. Reason for Adversarial RL

Motivated by the importance of zero-shot adaptation as discussed in Section II, we propose a RARL-based beam-tracking method. The key reason for using RARL is to develop the capability to overcome the training and test gap by: 1) regarding the gap as a disturbance from an adversarial agent that impedes the legitimate agent; 2) training both the adversarial and legitimate agent, thereby allowing the legitimate agent to experience more severe disturbances.

The explanation more specific to our beam-tracking problem is as follows: By training the adversarial agent to disturb the on-wire SBS with additional wind, the beam-tracking agent experiences more rapid displacements in the on-wire SBS than without such adversarial agents. We hypothesize that this well simulates a situation in which it would be difficult to correct the directional beams, where the actual wire mass or spring constant is smaller than that used for training. This means that the adversarial agent provides richer experiences to the beam-tracking agent in view of the existence of the training and test gap; hence, the beam-tracking agent would be expected to obtain a robust beam-tracking policy against these training and test gaps.

##### B. Overview of RARL-Based Beam-Tracking

The training procedure for the RARL-based beam-tracking is shown in Fig. 4. In the training scenario, as shown in Fig. 4, the protagonist corresponding to the beam-tracking agent learns to maximize the average received signal power. In contrast, the adversary learns to minimize the average received signal power by generating additional wind. To achieve these purposes, the protagonist and adversary observe a state, select an action, and observe a reward to update their NN from experienced transitions. In the test scenario shown in Fig. 5, the protagonist corrects the beam misalignment according to the policy learned in the training scenario. To examine the feasibility of zero-shot adaptation, the environmental parameters, for example, the spring constant  $k_0$  and total wire mass  $m$ , are varied between the training and test scenarios.

As an example of disturbance caused by an adversary, we assume that the adversary can affect the wind speed in the simulation. Thus, the adversary can append discontinuous additional wind to continuous wind in the environment. At every time step  $t$  in the training scenario, by considering the wind speed in the environment  $\mathbf{v}_e^{(k)} \in \mathbb{R}^3$  and the additional wind speed appended by the adversary  $\mathbf{v}_a^{(k)} \in \mathbb{R}^3$ , the wind speed  $\mathbf{v}_o^{(k)} \in \mathbb{R}^3$  in (2) is calculated as

$$\mathbf{v}_o^{(k)} = \mathbf{v}_e^{(k)} + \mathbf{v}_a^{(k)}. \quad (10)$$

Conversely, because the adversary does not exist in the test scenarios, the wind speed  $\mathbf{v}_o^{(k)} \in \mathbb{R}^3$  in (2) is given by

$$\mathbf{v}_o^{(k)} = \mathbf{v}_e^{(k)}. \quad (11)$$

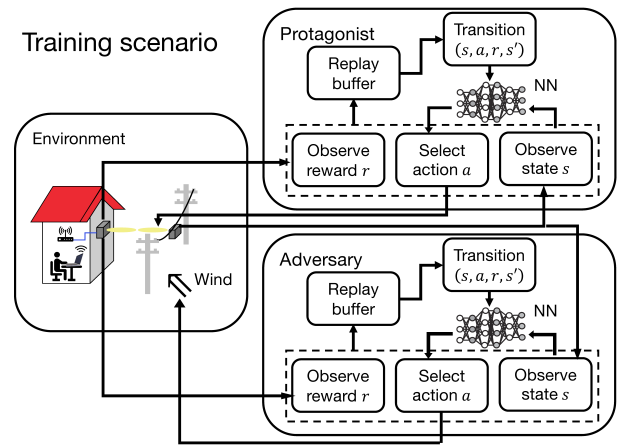


Fig. 4. Training scenario of the adversarial RL-based beam-tracking.

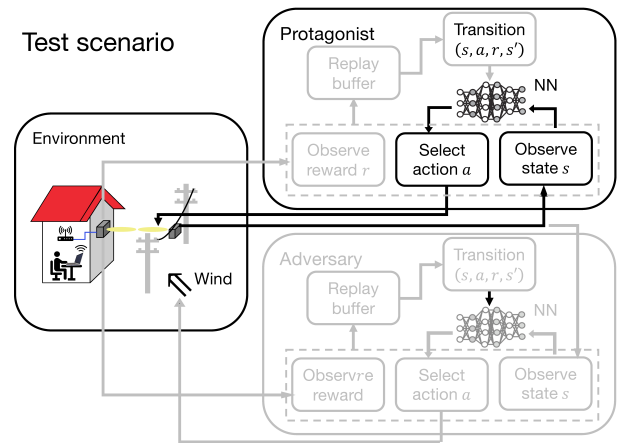


Fig. 5. Test scenario of the adversarial RL-based beam-tracking.

##### C. State, Action, and Reward

The state set  $\mathcal{S}$  of the protagonist and adversary is defined as

$$\mathcal{S} := \mathcal{S}_x \times \mathcal{S}_v \times \mathcal{S}_b, \quad (12)$$

where  $\mathcal{S}_x := \{\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^3\}$  and  $\mathcal{S}_v := \{\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^3\}$  denote the set of possible three-dimensional positions and velocities of the SBS, respectively. Note that for the practical implementation, these three-dimensional positions/velocities can be obtained if the SBS involves an accelerometer and then integrates the measured accelerations. This position/velocity acquisition involves an error, and this error may affect the accuracy of the beam-tracking. Nonetheless, we assume that these three-dimensional positions/velocities can be obtained without any measurement errors in view of the scope of this study. As discussed in Section I, our main scope is to validate the effectiveness of adding an adversary during training, thereby confirming the feasibility of our concept of zero-shot adaptation. This objective can be achieved by comparing the proposed method (with the adversary) with the baseline method (without the adversary), and the measurement error does not have a specific impact on this comparison because the inaccuracy of the beam-tracking due to the measurement error occurs commonly in both methods. Hence, this assumption is

sufficient to validate the contribution of this study, and delving into the measurement precisions of the positions/velocities is beyond the scope of this study.

In (12),  $\mathcal{S}_b := \{\mathbf{b} \mid \mathbf{b} \in \mathbb{R}^3\}$  denotes the set of possible beam directions, where beam direction at time step  $k$   $\mathbf{b}^{(k)}$  is given by

$$\mathbf{b}^{(k)} = [\sin \theta_s^{(k)} \cos \phi_s^{(k)}, \sin \theta_s^{(k)} \sin \phi_s^{(k)}, \cos \theta_s^{(k)}]^T. \quad (13)$$

These state settings are consistent with our previous works [20], [21] to ensure a fair comparison.

The action set of the protagonist  $\mathcal{A}_p$  is defined as

$$\mathcal{A}_p := \{\text{stay, up, down, left, right}\}, \quad (14)$$

where the action stay denotes that the beam direction is maintained. Moreover, the actions up, down, left, and right denote that the beam direction is moved up, down, left, and right by  $\beta$ , respectively. The actions for the zenith and azimuth angle  $a_\theta^{(k)}, a_\phi^{(k)}$  are given by

$$\begin{bmatrix} a_\theta^{(k)} \\ a_\phi^{(k)} \end{bmatrix} = \begin{cases} [0, 0]^T, & a_k = \text{stay}; \\ [-1, 0]^T, & a_k = \text{up}; \\ [1, 0]^T, & a_k = \text{down}; \\ [0, 1]^T, & a_k = \text{left}; \\ [0, -1]^T, & a_k = \text{right}. \end{cases} \quad (15)$$

The action set of the adversary  $\mathcal{A}_a$  is defined as

$$\mathcal{A}_a := \{\text{stay, up, down, left, right, front, back}\}, \quad (16)$$

where the action stay denotes that no additional wind is appended. Furthermore, in the training scenarios, the actions up, down, left, right, front, and back denote wind in the upward, downward, leftward, rightward, forward, and backward directions with a wind speed of  $v_a$ . The additional wind speed by the adversary  $\mathbf{v}_a^{(k)}$  is given by

$$\mathbf{v}_a^{(k)} = \begin{cases} [0, 0, 0]^T, & a_k = \text{stay}; \\ [0, 0, v_a]^T, & a_k = \text{up}; \\ [0, 0, -v_a]^T, & a_k = \text{down}; \\ [-v_a, 0, 0]^T, & a_k = \text{left}; \\ [v_a, 0, 0]^T, & a_k = \text{right}; \\ [0, v_a, 0]^T, & a_k = \text{front}; \\ [0, -v_a, 0]^T, & a_k = \text{back}. \end{cases} \quad (17)$$

The immediate reward is defined as the instantaneous received signal power at the next step, which is clipped using the technique in [31]:

$$r_{p,k} = \begin{cases} 1, & (P_r^{(k+1)} - b_c)/d_c > 1; \\ (P_r^{(k+1)} - b_c)/d_c, & -1 \leq (P_r^{(k+1)} - b_c)/d_c \leq 1; \\ -1, & (P_r^{(k+1)} - b_c)/d_c < -1, \end{cases} \quad (18)$$

where  $P_r^{(k+1)} := P_r(a^{(k+1)}, \theta^{(k+1)}, \phi^{(k+1)}, \theta_s^{(k+1)}, \phi_s^{(k+1)})$ . Moreover,  $b_c$  and  $d_c$  denote the offset and scale of the clipping, respectively. The immediate reward of the adversary  $r_{a,k}$  is defined by inverting the sign of that of the protagonist to encourage the adversary to disturb the beam-tracking agent, that is,  $r_{a,k} = -r_{p,k}$ .

#### D. Adversarial RL Algorithm

Given the definition of the state, action, and reward, the training of the protagonist beam-tracking agent and adversary is performed via deep Q-learning [31] for the two agents. In deep Q-learning, the following value, termed the optimal action-value function, is predicted via a neural network:

$$Q_{agent}^*(s, a) = \mathbb{E}_{\pi_{agent}^*} \left[ \sum_{k'=0}^{\infty} \gamma^{k'} r_{agent, k+1+k'} \mid s_k = s, a_{agent, k} = a \right], \quad s \in \mathcal{S}, a \in \mathcal{A}_{agent}, \quad (19)$$

where  $s_k$  and  $\gamma \in [0, 1)$  denote the state at the time step  $k$  and discount factor, respectively. For  $agent \in \{p, a\}$ , which indicates the protagonist or adversary,  $a_{agent, k}$  and  $r_{agent, k}$  denote the action and reward given to  $agent$  at time step  $k$ , respectively. The term  $\pi_{agent}^* : \mathcal{S} \rightarrow \mathcal{A}_{agent}$  is the optimal policy, which means the action rule to maximize the discounted reward. Although finding the optimal policy is the main objective of this algorithm, in deep Q-learning, we can find the optimal action value function first, and then we can determine the optimal policy by:  $\arg \max_{a \in \mathcal{A}_{agent}} Q_{agent}^*(s, a)$ . Hence,

the problem boils down to finding the optimal action-value function, which is conducted by training a neural network known as a deep Q network (DQN) in deep Q-learning such that the DQN is a good approximation of the optimal action-value function. Let  $Q_{agent}(s, a; \theta_{agent})$  denote the DQN for  $agent$ .

The procedure used to train the two DQNs is detailed below. Note that procedures 1 and 2 are conducted in every time step, whereas procedures 3, 4, 5 are performed on a per-episode basis. Here, we let the episode be the finite time steps for  $k = 1, 2, \dots, \lfloor T/\tau \rfloor$ .

**1. Exploring and Storing Experience.** In this procedure, the protagonist and adversary collect the ingredients to create training data, termed experience. The experience is defined as  $(s_k, a_{agent, k}, r_{agent, k}, s_{k+1})$  for  $agent$  and is collected while interacting with the environment, that is, taking the action, obtaining the rewards, and the subsequent states. In this procedure, the protagonist and adversary synchronously perform the action following  $\epsilon$ -greedy policies, which is a general assumption that includes the prior work for adversarial RL [22]. The experience for  $agent$  is stored in the experience memory denoted by  $\mathcal{D}_{agent}$ .

**2. Training DQNs.** Given the experience memories  $\mathcal{D}_{agent}$  for  $agent \in \{p, a\}$ , the DQNs are trained. In deep Q-learning, the DQN is trained to minimize the difference metric between  $Q_{agent}(s, a; \theta_{agent})$  and  $target_{agent} := r + \gamma \max_{a \in \mathcal{A}_{agent}} Q_{agent}(s, a; \theta_{agent}^-)$ , where  $Q_{agent}(s, a; \theta_{agent}^-)$  is termed the target network and is updated less frequently than the DQN. We refrain from delve into the details of the target network to enable us to focus on the training procedure. Interested readers are encouraged to refer to the paper of Mnih *et al.* [31].

The training procedure is as follows. First, we calculate  $target_{agent}$  by sampling the experience from  $\mathcal{D}_{agent}$  uniformly. Subsequently, we train each DQN to minimize the difference



metric between the DQN  $Q_{agent}(s, a; \theta_{agent})$  and  $target_{agent}$  via the Adam optimizer [32]. Note that the protagonist and adversary DQNs are trained synchronously, which is found to be sufficient to demonstrate the robustness of the protagonist against training-test gaps.

During this training procedure, we leveraged more advanced techniques in the evaluation described in Section V. More specifically, we leveraged the Huber loss [33] as the difference metric instead of taking the square of  $Q_{agent}(s, a; \theta_{agent}) - target_{agent}$ . In addition, we leveraged dueling DQN [34] in the experiment. However, detailing these techniques is beyond the objective of this section; hence, we detailed these techniques in the Appendix.

**3. Updating Target DQNs.** The parameters of the target DQNs  $\theta_{agent}^-$  are updated such that  $\theta_{agent}^-$  is equal to  $\theta_{agent}$ . This update is generally performed less frequently than that of the update the main DQNs [31]; hence, we conduct this procedure when every  $C$  episode elapses. This target network update is performed synchronously for the protagonist and adversary.

**4. Checking Protagonist Performance.** This procedure is performed to check whether an appropriate beam-tracking policy is learned regardless of the adversary. In this procedure, for additional *test* steps, the protagonist performs beam-tracking by greedily determining the action with respect to the DQN  $Q_p(s, a; \theta_p)$ , that is, it chooses the action by:  $\arg \max_{a \in \mathcal{A}_p} Q_p^*(s, a; \theta_p)$ , and the received power averaged for the steps is obtained. Therein, the adversary is not activated to check the protagonist performance in view of real deployments.

**5. Checking Adversary Performance.** This procedure is performed to check whether the adversary can with certainty learn a policy to disturb a protagonist. In this procedure, for additional *test* steps, the adversary disturbs the protagonist by greedily determining the action with respect to the DQN  $Q_a(s, a; \theta_a)$ . In contrast to the previous procedure, checking the performance of adversary still requires the existence of a protagonist, whereas using the protagonist in this learning procedure  $Q_p(s, a; \theta_p)$  underestimates the performance of the adversary because the protagonist may already be robust against the adversary. To avoid this, we prepared another proxy protagonist, *pre-trained* without the adversary, and it is this proxy protagonist that is disturbed by the adversary in this phase. This proxy protagonist did not learn a robust policy against the adversary; therefore, we can keep track of the performances of the adversary without underestimation. Note that by letting  $Q_p(s, a; \theta_{p,proxy})$  denote the DQN of the proxy protagonist, this protagonist also determines the action greedily with respect to  $Q_p(s, a; \theta_{p,proxy})$ . The performance of the adversary is measured by the average received power that is obtained in the proxy protagonist during this phase.

These procedures are iterated for a predefined number of episodes  $M$ . This value of  $M$  is determined to be sufficiently longer than the convergence of the protagonist performance, which could be measured in procedure 4. Note that during the training procedure, the adversary is always active except for procedure 4. The overall training procedure is summarized in

---

**Algorithm 1** Training the protagonist beam-tracking agent and adversary via adversarial RL

---

```

Initialize main DQN and target DQN of protagonist, i.e.,
 $Q_p(s, a; \theta_p)$  and  $Q_p(s, a; \theta_p^-)$ , respectively, and experience
memory of protagonist  $\mathcal{D}_p$ 
Initialize main DQN and target DQN of adversary, i.e.,
 $Q_a(s, a; \theta_a)$  and  $Q_a(s, a; \theta_a^-)$ , respectively, experience mem-
ory of adversary  $\mathcal{D}_a$ 
Observe initial state  $s_1$  by obtaining  $x_S^{(1)}$ , velocity
 $v_S^{(1)}$ , and beam direction  $b_S^{(1)}$ 
for Episode  $e = 1, 2, \dots, M$  do
  for time step  $k = 1, 2, \dots, \lfloor T/\tau \rfloor$  do
    1. Exploring and Storing Experience:
      Select actions of protagonist and adversary  $a_{p,k}$ 
      and  $a_{a,k}$  synchronously
      Calculate reward for protagonist  $r_{k,p}$  from (18)
      Calculate reward  $r_{k,a}$  by  $r_{k,a} \leftarrow -r_{k,p}$ 
      Observe state  $s_{k+1}$  by obtaining  $x_S^{(k+1)}$ , velocity
       $v_S^{(k+1)}$ , and beam direction  $b_S^{(k+1)}$ 
      Store experience  $(s_k, a_{p,k}, r_{p,k}, s_{k+1})$  into  $\mathcal{D}_p$ 
      Store experience  $(s_k, a_{a,k}, r_{a,k}, s_{k+1})$  into  $\mathcal{D}_a$ 
    2. Training DQNs:
      Synchronously update weights of main DQNs of pro-
      tagonist and adversary  $\theta_p, \theta_a$ 
    3. Updating Target DQNs:
      if  $e \equiv 0 \pmod{C}$  then
        Update target DQNs  $\theta_p^- \leftarrow \theta_p$  and  $\theta_a^- \leftarrow \theta_a$ 
      end if
    end for
    4. Checking Protagonist Performance:
      for time step  $k = \lfloor T/\tau \rfloor + 1, \dots, \lfloor T/\tau \rfloor + test$  do
        Select actions of protagonist  $a_{p,k}$  greedily w.r.t.
         $Q_p(s, a; \theta_p)$ 
        Observe state  $s_{k+1}$  in the same way as that in
        procedure 1
        Observe received power
      end for
    5. Checking Adversary Performance:
      for time step  $k = \lfloor T/\tau \rfloor + 1, \dots, \lfloor T/\tau \rfloor + test$  do
        Select actions of proxy protagonist  $a_{p,k}$  greed-
        ily w.r.t.  $Q_p(s, a; \theta_{p,proxy})$  (pre-trained without
        adversary)
        Select actions of adversary  $a_{a,k}$  greedily w.r.t.
         $Q_a(s, a; \theta_a)$ 
        Observe state  $s_{k+1}$  in the same way as that in
        procedure 1
        Observe received power
      end for
      Check average received power
    end for
  
```

---

Algorithm 1.

## V. SIMULATION RESULTS

### A. Simulation Parameters

The simulation parameters are listed in Table III, where  $\mathbf{E} \in \mathbb{R}^{3 \times 3}$  is an identity matrix. The SBS is installed at the midpoint of the overhead messenger wire, which experiences



Fig. 6. Antenna pattern. The vertical angle  $\theta$  was fixed at  $90^\circ$ .

Fig. 7. Architecture of the NN of the protagonist and adversary.

significant movement as a result of the wind. The wind speed in the environment  $v_e^{(k)}$  is given by

$$v_e^{(k)} = \left[ 5 \sin \frac{2\pi k\tau}{4}, 5 \sin \frac{2\pi k\tau}{6}, 5 \sin \frac{2\pi k\tau}{8} \right]^T, \quad (20)$$

which follows our previous studies [20], [21].

1) *Antenna Pattern*: Fig. 6 shows the antenna pattern based on the simulation parameters listed in Table III. The red trace represents the transmission antenna gain with the array factor  $AF(\theta, \phi, \theta_s, \phi_s)$ . The blue curve represents the transmission antenna gain without the array factor, that is,  $AF(\theta, \phi, \theta_s, \phi_s) = 0$  in (4). The directivity of the transmission antenna gain with the array factor is higher than that of the transmission antenna gain without the array factor. In the transmission antenna gain with the array factor,  $3.6^\circ$  is the positive minimum value of the local minimum values.

2) *Architecture of the Neural Network*: We used a neural network with four hidden layers, as shown in Fig. 7, where  $|\mathcal{A}|$  denotes the number of actions. The number of actions of the protagonist  $|\mathcal{A}_p| = 5$  and the number of actions of the adversary  $|\mathcal{A}_a| = 7$ . The hidden layers were all fully connected and had 32 units. The activation function of the hidden layers is the rectified linear unit  $R(x)$ , which is given by

$$R(x) = \max\{x, 0\}. \quad (21)$$

Adam [32] was used as the gradient descent method, and the learning rates of the protagonist and adversary  $\alpha_p, \alpha_a$  were  $10^{-3}$ .

TABLE III  
SIMULATION PARAMETERS

|  |                                |
|--|--------------------------------|
| Height of endpoints of wire $h_w$              | 5 m                            |
| Distance between endpoints $d_w$               | 10 m                           |
| Height of gateway BS $h_r$                     | 5 m                            |
| Distance between wire and gateway BS $d_r$     | 5 m                            |
| Transmission power $P_t$                       | 23 dBm                         |
| Radio-wave wavelength $\lambda$                | 5 mm                           |
| Receiver antenna gain $A_r$                    | 8 dBi                          |
| Gravitational acceleration $g$                 | $[0, 0, -9.8] \text{ ms}^{-2}$ |
| Spring constant $k_0$                          | $100 \text{ Nm}^{-1}$          |
| Drag constant $c_0$                            | $1 \text{ s}^{-1}$             |
| Number of points $N$                           | 11                             |
| Total wire mass $m$                            | 10 kg                          |
| Covariance matrix of the wind speed $V_0$      | $0.1E$                         |
| Vertical 3 dB beamwidth $\theta_{3\text{dB}}$  | $65^\circ$                     |
| Horizontal 3 dB beamwidth $\phi_{3\text{dB}}$  | $65^\circ$                     |
| Side-lobe level limit $SLA_v$                  | 30 dB                          |
| Front-back ratio $A_m$                         | 30 dB                          |
| Number of vertical elements $n_v$              | 32                             |
| Number of horizontal elements $n_H$            | 32                             |
| Zenith steering angle $\theta_s$               | $90^\circ$                     |
| Azimuth steering angle $\phi_s$                | $0^\circ$                      |
| Vertical spacing distance $\Delta_v$           | 2.5 mm                         |
| Horizontal spacing distance $\Delta_H$         | 2.5 mm                         |
| Number of episodes $M$                         | 400                            |
| Point of SBS installation                      | Point 6                        |
| Observation time $T$                           | 10 s                           |
| Observation interval $\tau$                    | 0.01 s                         |
| Angle when moving beam direction $\beta$       | $1^\circ$                      |
| Constant in (18) $b_c$                         | 27                             |
| Constant in (18) $d_c$                         | 3                              |
| Exploration rate for $\epsilon$ -greedy policy | 0.2                            |
| Discount factor $\gamma$                       | 0.99                           |
| Target DQN frequency $C$                       | 5                              |
| Wind speed added by the adversary $v_a$        | 10 m/s                         |

### B. Baseline Method

To evaluate the average received signal power of the learned policy of the protagonist, we compared the proposed method with the *stay method* and the *upper-limit method*. Because these compared methods aim to evaluate the policy of the protagonist without the disturbance caused by the adversary, the adversary does not exist in the stay or upper-limit methods. In the stay method, the beam direction is fixed in the initial beam direction. The policy  $\pi(a | s)$  of the stay method is given by

$$\pi(a | s) = \begin{cases} 1, & a = \text{stay}; \\ 0, & a = \text{up, down, left, right}. \end{cases} \quad (22)$$

In the upper-limit method, the actions maximize the transmission antenna gain such that the superior performance is delivered in terms of the received signal power.

To evaluate the robustness of the learned policy of the protagonist, we compared the proposed method with the *no adversary method* and *random 10 m/s method*. In the no adversary method, the policy of the protagonist is learned in the scenario in which the adversary appends no additional wind, i.e.,  $v_a = 0$ . In the random 10 m/s method, the policy of the protagonist is learned in the scenario in which the adversary takes each action at random with equal probability.




Fig. 8. Learning curve of the protagonist and the adversary. Only in the adversary performance in the blue curve, the lower received power is better. As the episode elapses, the performance of both the protagonist and adversary increases as the episode elapses.

### C. Learning Curve

The result of the performance check of the protagonist and adversary during training is shown in Fig. 8, along with the performance of the baseline methods. The red curve shows the performance of the protagonist obtained during the procedure “4. Checking Protagonist Performance” whereas the blue curve shows the performance of the adversary obtained during the procedure “5. Checking Adversary Performance,” which is detailed in Section IV-D. Note that only with respect to the performance of the adversary (represented by the blue curve), lower received power is better because the performance of the adversary should be measured with the extent to which it disturbs the proxy protagonist. As shown in Fig. 8, regardless of the adversarial training, the performance of both the protagonist and adversary increases as the episode elapses. This first shows the feasibility of the protagonist achieving an appropriate beam-tracking policy that converges closely to the upper limit regardless of the adversary. Moreover, although the received power fluctuates, the adversary can basically obtain the policy of disturbing the protagonist as the episode elapses.

### D. Robustness of the Learned Policy

To demonstrate the robustness to variations in the spring constant, we compared the average received signal power of the proposed method with that of the no adversary and random 10 m/s methods, as shown in Fig. 9. In Fig. 9, the spring constant of the messenger wire is shown on the horizontal axis, where the training parameter of the spring constant is represented by the dashed line. When the spring constant  $k_0 = 100 \text{ Nm}^{-1}$ , which is the setting in the training, all the methods, including baseline methods, exhibited an almost identical amount of received power, that is, approximately  $-12.9 \text{ dBm}$ . However, in the scenario of when the spring constant was lower than that in training, that is,  $k_0 = 10 \text{ Nm}^{-1}$ , the average received signal power of the proposed method was  $-13.2 \text{ dBm}$ , whereas that of the no adversary and random 10 m/s methods dropped to  $-14.5$  and  $-13.8 \text{ dBm}$ , respectively. Indeed, the random 10 m/s method (orange curve) is advantageous to enable the protagonist to learn a robust policy

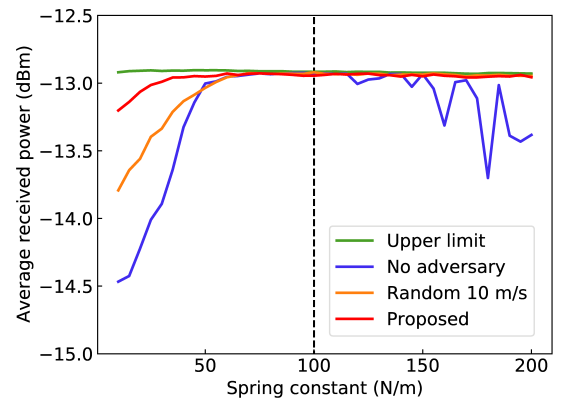


Fig. 9. Robustness to variations in the spring constant of the overhead messenger wire. In the test scenarios, the average received signal power of the proposed method is more robust to different spring constants than the compared methods. The dashed line represents the training parameter of the spring constant.




Fig. 10. Robustness to variations in total wire mass. In the test scenarios, the average received signal power of the proposed method is more robust to various total wire masses than the compared methods. The dashed line represents the training parameter of the total wire mass.

when compared to no adversary method. However, in contrast to the random 10 m/s method, the proposed method with the adversary is more effective for the protagonist to learn a robust policy.

To evaluate the robustness to the variations in the total wire mass, we compared the average received signal power of the proposed method with that of the no adversary and random 10 m/s methods, as shown in Fig. 10. In Fig. 10, the total wire mass is shown on the horizontal axis, where the training parameter of the total wire mass is represented by the dashed line, that is, 10 kg. Similar to the results pertaining to the spring constant, the proposed methods exhibited a larger amount of received power for various amounts of the total wire mass. Hence, the proposed method achieves more robust beam-tracking against the training-and-test gaps in terms of the total wire mass. Note that these achievements are without any adaptive fine-tuning of the total wire mass and hence demonstrates the feasibility of the aforementioned zero-shot adaptation in mmWave beam-tracking.

We discuss the robustness to variations in the spring constants and total wire mass by plotting the average received signal power for various values of the spring constant and total

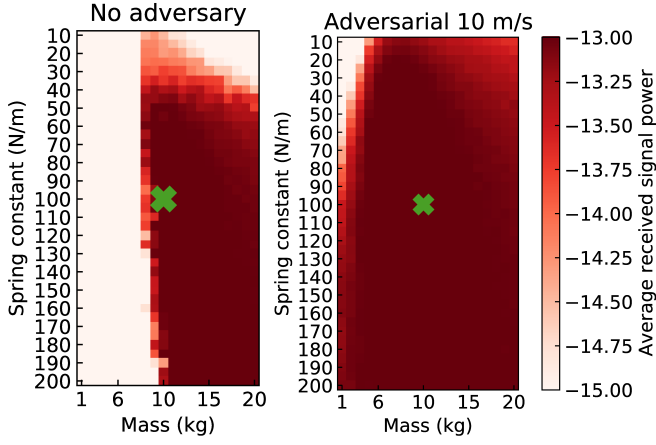


Fig. 11. Heat map depicting the robustness to variations in the spring constant and the total wire mass of the messenger wire. In the test scenarios, the average received signal power of the proposed method is more robust to the variations in the total wire mass than the no adversary method. The cross mark represents the training parameters of the spring constant and total wire mass.

wire mass in Fig. 11. The color of the heat map represents the average received signal power, where the cross mark represents the training parameters of the spring constant and total wire mass. In Fig. 11, the area in which the average amount of received power is higher, is shown in red, whereas that ohin which the average received power is low, is shown in white. Fig. 11 demonstrates that the area representing the high average received power of the proposed method is wider than that of the no adversary method. This implies that the proposed method enables the adaptation of various ranges of test settings for beam-tracking without the requirement for adaptive fine-tuning; thus, in this sense, this result demonstrates the feasibility of zero-shot adaptation in learning-based beam-tracking.

## VI. CONCLUSION

We discussed zero-shot adaptation in learning-based mmWave beam-tracking. To demonstrate the feasibility of zero-shot adaptation, we proposed an adversarial RL-based beam-tracking method to obtain a robust beam-tracking policy to overcome the differences between the training and test scenarios, such as variations in the wire tension and the total wire mass. We developed an RARL-based algorithm in which the adversary generated additional wind. We demonstrated that the proposed method is more robust, not only than the no adversary method but also than the random 10 m/s method, when assuming an adversary without training. This shows that disturbance by an intelligent adversary increased the robustness of the beam-tracking policy to variations in the wire tension, that is, the spring constant or total wire mass. The improved robustness was greated than for an adversary without training.

The insights provided in this paper open numerous future challenges in learning-based systems in mmWave networks. As noted in Section I-B, the training-test gap in terms of environmental parameters affecting node dynamics may be problematic not only in beam-tracking in non-mobile scenarios, but also in a handover problem in mobile scenarios,

whereas design policies for handover are different from those in beam-tracking discussed in this study. For example, in contrast with the beam-tracking problem in this study, which considered one metric, a handover problem in mobile scenarios should jointly consider multiple metrics, e.g., the received powers and the frequency of handover operations. Extending the concept of zero-shot adaptation provided in Section II to such multi-metric problems and designing the adversarial agent that enables protagonist agents to achieve satisfactory performances are included in a future challenge of this study.

## APPENDIX

**Huber Loss [33].** The difference between  $Q_{agent}(s, a; \theta_{agent})$  is measured by the Huber loss [33], which is given by:

$$\mathcal{L}(x) = \begin{cases} x^2/2, & x \leq 1; \\ |x| - 0.5, & \text{otherwise,} \end{cases} \quad (23)$$

where

$$x = r + \gamma \max_{a'} Q(s', a'; \theta^-) - Q(s, a; \theta_i). \quad (24)$$

**Dueling DQN (DDQN) [34].** In dueling DQN, each Q network is divided into two networks. More specifically, the action-value function  $Q^\pi(s, a)$  is divided into a state value function  $V^\pi(s)$  and an advantage function  $A^\pi(s, a)$ . The state-value and advantage functions are expressed as follows:

$$\begin{aligned} V^\pi(s) &= \mathbb{E}_{a \sim \pi(s)}[Q^\pi(s, a)], \\ A^\pi(s, a) &= Q^\pi(s, a) - V^\pi(s). \end{aligned} \quad (25)$$

By using the state-value function approximated by a neural network  $V(s; \theta_V)$  and the advantage function approximated by NN  $A(s, a; \theta_A)$ , the Q-network  $Q(s, a; \theta^-)$  is expressed as follows:

$$\begin{aligned} Q(s, a; \theta) &= V(s; \theta_V) \\ &+ \left( A(s, a; \theta_A) - \frac{1}{|\mathcal{A}|} \sum_{a'} A(s, a'; \theta_A) \right), \end{aligned} \quad (26)$$

where  $\theta_V$  and  $\theta_A$  are the weights used by neural network for approximating the state value and advantage functions, respectively. Note that  $\theta = [\theta_V, \theta_A]$ .

## REFERENCES

- [1] Y. Niu, Y. Li, D. Jin, L. Su, and A. V. Vasilakos, "A survey of millimeter wave communications (mmWave) for 5G: Opportunities and challenges," *Wireless netw.*, vol. 21, no. 8, pp. 2657–2676, Nov. 2015.
- [2] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-wave cellular wireless networks: Potentials and challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, Mar. 2014.
- [3] Q. Cao, H. Rutagemwa, F. Zhou, P. Yu, L. Feng, W. Li, A. Xiong, and X. Qiu, "Capacity enhancement for mmWave multi-beam satellite-terrestrial backhaul via beam sharing," in *Proc. IEEE ICC*. Kansas City, MO, USA: IEEE, May 2018, pp. 1–6.
- [4] L. Wei, R. Q. Hu, Y. Qian, and G. Wu, "Key elements to enable millimeter wave communications for 5G wireless systems," *IEEE Wireless Commun.*, vol. 21, no. 6, pp. 136–143, Dec. 2014.
- [5] "Wireless LAN medium access control (MAC) and physical layer (PHY) specification," IEEE Std. 802.11ad-2012.
- [6] G. C. Alexandropoulos, "Position aided beam alignment for millimeter wave backhaul systems with large phased arrays," in *Proc. IEEE CAMSAP*, Curacao, Netherlands Antilles, Dec. 2017, pp. 1–5.



- [7] C. Zhang, D. Guo, and P. Fan, "Tracking angles of departure and arrival in a mobile millimeter wave channel," in *Proc. IEEE ICC*, Kuala Lumpur, Malaysia, May 2016, pp. 1–6.
- [8] V. Va, H. Vikalo, and R. W. Heath, "Beam tracking for mobile millimeter wave communication systems," in *Proc. IEEE GlobalSIP*, Washington, WA, USA, Dec. 2016, pp. 743–747.
- [9] S. Shaham, M. Ding, M. Kokshoorn, Z. Lin, S. Dang, and R. Abbas, "Fast channel estimation and beam tracking for millimeter wave vehicular communications," *IEEE Access*, vol. 7, pp. 141 104–141 118, Sep. 2019.
- [10] S. G. Larew and D. J. Love, "Adaptive beam tracking with the unscented kalman filter for millimeter wave communication," *IEEE Signal Process. Lett.*, vol. 26, no. 11, pp. 1658–1662, Sep. 2019.
- [11] J. Lim, H.-M. Park, and D. Hong, "Beam tracking under highly nonlinear mobile millimeter-wave channel," *IEEE Commun. Lett.*, vol. 23, no. 3, pp. 450–453, Mar. 2019.
- [12] T. Lin and Y. Zhu, "Beamforming design for large-scale antenna arrays using deep learning," *IEEE Wireless Commun. Lett.*, vol. 9, no. 1, pp. 103–107, Sep. 2019.
- [13] A. M. Elbir and K. V. Mishra, "Joint antenna selection and hybrid beamformer design using unquantized and quantized deep learning networks," *IEEE Trans. Wireless Commun.*, vol. 19, no. 3, pp. 1677–1688, Dec. 2019.
- [14] —, "Deep learning design for joint antenna selection and hybrid beamforming in massive MIMO," in *Proc. IEEE ISAP*, Atlanta, GA, USA, Jul. 2019, pp. 1585–1586.
- [15] Y. Wang, M. Narasimha, and R. W. Heath, "Mmwave beam prediction with situational awareness: A machine learning approach," in *Proc. IEEE SPAWC*, Kalamata, Greece, Jun. 2018, pp. 1–5.
- [16] Y. Wang, A. Klautau, M. Ribero, A. C. Soong, and R. W. Heath, "Mmwave vehicular beam selection with situational awareness using machine learning," *IEEE Access*, vol. 7, pp. 87 479–87 493, Jun. 2019.
- [17] A. Klautau, P. Batista, N. González-Prelcic, Y. Wang, and R. W. Heath, "5G MIMO data for machine learning: Application to beam-selection using deep learning," in *Proc. IEEE ITA*, San Diego, CA, USA, Feb. 2018, pp. 1–9.
- [18] R. Wang, O. Onireti, L. Zhang, M. A. Imran, G. Ren, J. Qiu, and T. Tian, "Reinforcement learning method for beam management in millimeter-wave networks," in *Proc. IEEE UCET*, Glasgow, United Kingdom, Aug. 2019, pp. 1–4.
- [19] F. B. Mismar, B. L. Evans, and A. Alkhateeb, "Deep reinforcement learning for 5G networks: Joint beamforming, power control, and interference coordination," *IEEE Trans. Commun.*, vol. 68, no. 3, pp. 1581–1592, Dec. 2019.
- [20] M. Shinzaki, Y. Koda, K. Yamamoto, T. Nishio, M. Morikura, H. Chunhsiang, Y. Shirato, and N. Kita, "Deep reinforcement learning-based beam tracking from mmwave antennas installed on overhead messenger wires," in *Proc. IEEE VTC-Fall*, Held Online, Dec. 2020, pp. 1–6.
- [21] Y. Koda, M. Shinzaki, K. Yamamoto, T. Nishio, M. Morikura, Y. Shirato, D. Uchida, and N. Kita, "Millimeter wave communications on overhead messenger wire: Deep reinforcement learning-based predictive beam tracking," *IEEE Trans. Cogn. Commun. Netw. (early access)*, Apr. 2021.
- [22] L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, "Robust adversarial reinforcement learning," in *Proc. ICML*, Sydney, Australia, Aug. 2017, pp. 2817–2826.
- [23] W. Wang, V. W. Zheng, H. Yu, and C. Miao, "A survey of zero-shot learning: Settings, methods, and applications," *ACM Trans. Intell. Syst. and Technol.*, vol. 10, no. 2, pp. 1–37, Feb. 2019.
- [24] G. M. Gladwell, "Inverse problems in vibration," *Appl. Mech. Rev.*, vol. 39, no. 7, pp. 1013–1018, Jul. 1986.
- [25] R. Zárate-Minano, F. M. Mele, and F. Milano, "SDE-based wind speed models with weibull distribution and exponential autocorrelation," in *Proc. IEEE PESGM*, Boston, MA, USA, Jul. 2016, pp. 1–5.
- [26] M. Rebato, L. Resteghini, C. Mazzucco, and M. Zorzi, "Study of realistic antenna patterns in 5G mmWave cellular scenarios," in *Proc. IEEE ICC*, Kansas City, USA, May 2018, pp. 1–6.
- [27] H. Shiri, J. Park, and M. Bennis, "Massive autonomous UAV path planning: A neural network based mean-field game theoretic approach," in *Proc. IEEE GLOBECOM*, Honolulu, HI, USA, Dec. 2019, pp. 1–6.
- [28] H. T. Friis, "A note on a simple transmission formula," *Proceedings of the IRE*, vol. 34, no. 5, pp. 254–256, May 1946.
- [29] S. Geng, J. Kivinen, X. Zhao, and P. Vainikainen, "Millimeter-wave propagation channel characterization for short-range wireless communications," *IEEE Trans. Veh. Technol.*, vol. 58, no. 1, pp. 3–13, Jan. 2008.
- [30] 3GPP TR 37.840 v12.1.0, "Technical specification group radio access network; study of radio frequency (RF) and electromagnetic compatibility (EMC) requirements for active antenna array system (AAS) base station," Tech. Rep., Dec. 2013.
- [31] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, San Diego, CA, USA, May 2015, pp. 1–15.
- [33] D. Varga, D. Saupe, and T. Szirányi, "DeepRN: A content preserving deep architecture for blind image quality assessment," in *Proc. IEEE ICME*, San Diego, CA, USA, Jul. 2018, pp. 1–6.
- [34] Z. Wang, T. Schaul, M. Hessel, H. Van Hasselt, M. Lanctot, and N. De Freitas, "Dueling network architectures for deep reinforcement learning," in *Proc. ICML*, New York City, NY, USA, Jun. 2016, pp. 1995–2003.



**Masao Shinzaki** received the B.E. degree in electrical and electronic engineering from Kyoto University in 2019. He is currently studying toward the M.E. degree at the Graduate School of Informatics from Kyoto University. He is a member of the IEICE.



**Yusuke Koda** received the B.E. degree in electrical and electronic engineering from Kyoto University in 2016 and the M.E. and Ph.D. degrees (Informatics) at the Graduate School of Informatics from Kyoto University in 2018 and 2021, respectively. He is currently a postdoctoral researcher at the Centre for Wireless Communications, University of Oulu. In 2019, he visited Centre for Wireless Communications, University of Oulu, Finland to conduct collaborative research. He received the VTS Japan Young Researcher's Encouragement Award in 2017 and TELECOM System Technology Award in 2020. He was a Recipient of the Nokia Foundation Centennial Scholarship in 2019.



**Koji Yamamoto** (S'03–M'06–SM'20) received the B.E. degree in electrical and electronic engineering from Kyoto University in 2002, and the master and Ph.D. degrees in Informatics from Kyoto University in 2004 and 2005, respectively. From 2004 to 2005, he was a research fellow of the Japan Society for the Promotion of Science (JSPS). Since 2005, he has been with the Graduate School of Informatics, Kyoto University, where he is currently an associate professor. From 2008 to 2009, he was a visiting researcher at Wireless@KTH, Royal Institute of Technology (KTH) in Sweden. He serves as an editor of IEEE Wireless Communications Letters and Journal of Communications and Information Networks, a track co-chair of APCC 2017, CCNC 2018, APCC 2018, and CCNC 2019, and a vice co-chair of IEEE ComSoc APB CCC. He was a tutorial lecturer in ICC 2019. His research interests include radio resource management, game theory, and machine learning. He received the PIMRC 2004 Best Student Paper Award in 2004, the Ericsson Young Scientist Award in 2006. He also received the Young Researcher's Award, the Paper Award, SUEMATSU-Yasuharu Award from the IEICE of Japan in 2008, 2011, and 2016, respectively, and IEEE Kansai Section GOLD Award in 2012. He is a senior member of the IEICE and a member of the Operations Research Society of Japan.



**Takayuki Nishio** (S'11-M'14-SM'20) received the B.E. degree in electrical and electronic engineering and the master's and Ph.D. degrees in informatics from Kyoto University in 2010, 2012, and 2013, respectively. He was an assistant professor in communications and computer engineering with the Graduate School of Informatics, Kyoto University from 2013 to 2020. He is currently an associate professor at the School of Engineering, Tokyo Institute of Technology, Japan. From 2016 to 2017, he was a visiting researcher in Wireless Information Network Laboratory (WINLAB), Rutgers University, United States. His current research interests include machine learning-based network control, machine learning in wireless networks, and heterogeneous resource management.

work Laboratory (WINLAB), Rutgers University, United States. His current research interests include machine learning-based network control, machine learning in wireless networks, and heterogeneous resource management.



**Masahiro Morikura** received B.E., M.E. and Ph.D. degree in electronic engineering from Kyoto University, Kyoto, Japan in 1979, 1981 and 1991, respectively. He joined NTT in 1981, where he was engaged in the research and development of TDMA equipment for satellite communications. From 1988 to 1989, he was with the communications Research Centre, Canada as a guest scientist. From 1997 to 2002, he was active in standardization of the IEEE802.11a based wireless LAN. He received Paper Award, Achievement Award and Distinguished

Achievement and Contributions Award from the IEICE in 2000, 2006 and 2019, respectively. He also received Education, Culture, Sports, Science and Technology Minister Award in 2007 and Maejima Award from the Teishin association in 2008 and the Medal of Honor with Purple Ribbon from Japan's Cabinet Office in 2015. Dr. Morikura is now a professor of the Graduate School of Informatics, Kyoto University. He is a Fellow of the IEICE and a member of IEEE.



**Yushi Shirato** received B.E., M.E and D.E. degrees in electrical engineering from Tokyo University of Science in 1990, 1992 and 2018, respectively. Since joining NTT Wireless Systems Laboratories in 1992, he has been engaged in R&D of adaptive equalizers, modems for fixed wireless access systems, and software defined radio systems. He is currently engaged in R&D of millimeter-wave-band very high throughput fixed wireless backhaul systems. He received the Best Paper Award in 2018 and the Young Engineer's Award in 2000 from IEICE, and the 18th Telecom System Technology Award from the Telecommunications Advancement Foundation in 2003. He is a senior member of IEICE.

com System Technology Award from the Telecommunications Advancement Foundation in 2003. He is a senior member of IEICE.



**Daisei Uchida** received B.E. and M.E. degrees in applied physics from Tokyo Institute of Technology in 1994 and 1997, respectively. Since joining NTT Wireless Systems Laboratories in 1997, he has been engaged in R&D of the access scheme for satellite systems, radio channel assignment and tree topology of PHS-based local positioning and information systems to form wireless multi-hop networks autonomously, radio channel assignment and transmission power control for mesh-type broadband fixed wireless access systems, modulation-demodulation

schemes for MIMO-OFDM systems, and modulation-demodulation schemes and radio circuit designs for Low Power Wide Area systems. He is currently engaged in R&D of distributed antennas systems using high frequency bands for 6G generation mobile technologies. He is a member of the Institute of Electronics, Information and Communication Engineers (IEICE) of Japan. He received the Young Engineers Award from IEICE in 2001.



**Naoki Kita** received his B.E. degree from Tokyo Metropolitan Institute of Technology, Tokyo, Japan in 1994, and received his M.E. and Ph.D. degrees from Tokyo Institute of Technology, Tokyo, Japan in 1996 and 2007, respectively. Since joining NTT in 1996, he has been engaged in the research of propagation characteristics for wireless access systems, the development of future satellite communication systems, and international standardization on radio wave propagation. From 2009 to 2010, he was a visiting scholar at Stanford University, CA, USA.

From 2013 to 2014, he was a visiting research scholar at Waseda University, Tokyo, Japan. He is currently a Senior Research Engineer, Supervisor, Group Leader at NTT Access Network Service Systems Laboratories, where he engages in the research and development of future wireless access network systems. He received the IEICE Young Researcher's Award, the IEICE Communications Society Best Paper Award, and IEICE Best Paper Award in 2002, 2010, and 2014, respectively. He also received the Best Paper Award in International Symposium on Antennas and Propagation 2016 (ISAP2016) in 2016. He is a senior member of IEICE and a member of IEEE.