

Spatially Consistent Air-to-Ground Channel Modeling via Generative Neural Networks

Amedeo Giuliani¹, Rasoul Nikbakht², Giovanni Geraci³, *Senior Member, IEEE*,
Seongjoon Kang⁴, *Graduate Student Member, IEEE*, Angel Lozano⁵, *Fellow, IEEE*,
and Sundeep Rangan⁶, *Fellow, IEEE*

Abstract—This letter proposes a generative neural network architecture for spatially consistent air-to-ground channel modeling. The approach considers the trajectories of uncrewed aerial vehicles along typical urban paths, capturing spatial dependencies within received signal strength (RSS) sequences from multiple cellular base stations (gNBs). Through the incorporation of conditioning data, the model accurately discriminates between gNBs and drives the correlation matrix distance between real and generated sequences to minimal values. This enables evaluating performance and mobility management metrics with spatially (and by extension temporally) consistent RSS values, rather than independent snapshots. For some tasks underpinned by these metrics, say handovers, consistency is essential.

Index Terms—Cellular network, channel model, drone, 5G, generative neural network, ray tracing, uncrewed aerial vehicle (UAV).

I. INTRODUCTION

NEXT-GENERATION mobile networks are envisioned to reliably connect uncrewed aerial vehicles (UAVs) [1], [2], [3], [4]. This will require a re-engineering of existing deployments, e.g., via multiantenna techniques, dedicated infrastructure, and cellular-satellite integration [5], [6], [7]. Accurate channel models are crucial for evaluating these and other solutions.

Extending statistical channel models to air-to-ground scenarios is challenging due to the complex dependencies on UAV altitude, orientation, and building height, among other aspects [8], [9], [10]. Data-driven approaches have been put forth for site-specific channel modeling, mapping spatial locations to channel parameters through regression [11], [12], [13]. Generative neural networks (GNNs) provide an alternative for non-site-specific propagation

modeling [14], [15], [16]. However, existing works can only produce independent channel snapshots, failing to capture how signals fluctuate during motion. While adequate for performance evaluations where a marginal distribution suffices, say for coverage, this is insufficient and might cause artifacts when designing networks to cope with UAV mobility.

This letter proposes a new GNN architecture for spatially consistent air-to-ground channel modeling. Specifically, the large-scale channel behavior is represented, as embodied by the local-average received signal strength (RSS); this subsumes every aspect save for the small-scale fading. In fact, the transmit power is set to 0 dBm and the antennas are taken to be omnidirectional, such that the RSS (in dBm) exactly equals the large-scale channel gain (in dB). The approach considers a UAV flying along random typical trajectories in an urban area. To capture dependencies within sequences RSS from multiple base stations (gNBs) on buildings of varying heights, a generative adversarial network (GAN) is introduced that incorporates two types of conditioning data: the distance sequence from a specific gNB, and the gNB index. This allows generating RSS sequences along designated trajectories corresponding to a particular gNB. Ray tracing is used for data provisioning, although the approach is compatible with measured data. This model:

- Successfully learns the marginal distribution of RSS from different gNBs and accurately discriminates between them, despite their similarity.
- Generates spatially consistent RSS sequences, with the correlation matrix distance between real and generated sequences converging to minimal values. Data augmentation through sequence self-convolution further enhances the accuracy of the model.

Once trained, the model can serve as the workhorse of system-level evaluations consisting of:

- Producing UAV trajectories and gNB locations stochastically, based on a deployment model providing conditioning data for each UAV-gNB link.
- Sampling random vectors from a prior distribution and feeding them to the model, along with the conditioning data, thus obtaining a sequence of RSS values.

With a faster spatial sampling rate, a similar approach could be employed to incorporate the small-scale fading and produce the full multipath channel response, including path gains, delays, and angles.

II. PROBLEM FORMULATION

A UAV and several gNBs are considered, respectively acting as receiver and transmitters. By reciprocity, their roles are

Manuscript received 26 November 2023; revised 12 January 2024; accepted 4 February 2024. Date of publication 7 February 2024; date of current version 11 April 2024. This work was supported in part by the Spanish Research Agency under Grant PID2021-123999OB-I00 and Grant CEX2021-001195-M; in part by the European Union NextGenerationEU through the UNICO 5G I+D SORUS Project; in part by the UPF-Fractus Chair; in part by ICREA; in part by NSF under Grant 1952180, Grant 2133662, Grant 2236097, Grant 2148293, and Grant 1925079; in part by NTIA; and in part by the Industrial Affiliates of NYU WIRELESS. The associate editor coordinating the review of this article and approving it for publication was G. Zheng. (Corresponding author: Giovanni Geraci.)

Amedeo Giuliani and Rasoul Nikbakht are with the Centre Tecnològic de Telecomunicacions de Catalunya, 08860 Castelldefels, Spain.

Giovanni Geraci is with Telefónica, 28050 Madrid, Spain, and also with the Department of Information and Communication Technologies, Universitat Pompeu Fabra, 08018 Barcelona, Spain (e-mail: giovanni.geraci@upf.edu).

Seongjoon Kang and Sundeep Rangan are with NYU Tandon School of Engineering, Brooklyn, NY 11201 USA.

Angel Lozano is with the Department of Information and Communication Technologies, Universitat Pompeu Fabra, 08018 Barcelona, Spain.

Digital Object Identifier 10.1109/LWC.2024.3363653




Fig. 1. Simulation area showing an example of a UAV trajectory (red) between buildings (green) and the locations of three gNBs (Δ).

interchangeable. On a circumscribed area in the city of Boston, the gNBs located on buildings of varying heights and the UAV follows random typical trajectories at a fixed height of 30 m. Fig. 1 depicts a typical such trajectory.

Spatially Consistent Generative Model: Each gNB-UAV link is characterized by the RSS. The collection of RSS values for the k th trajectory are denoted by

$$\mathbf{x}_i^k = [p_1, \dots, p_N], \quad (1)$$

where i is the gNB identifier, N the number of spatial steps, and p_n the RSS value at the n th step. Similarly, the evolution of the UAV-gNB distance is denoted by

$$\mathbf{u}_i^k = [d_1, \dots, d_N], \quad (2)$$

where d_n is the 3D distance at the n th step. In the sequel, we employ a conditioning variable c equal to the gNB identifier i , yet the methodology can be extended to more general conditioning variables, such as the height and gNB type. The goal is to capture dependencies within the RSS sequences for multiple specific gNBs across a set of typical trajectories, i.e., to model the conditional distribution $p(\mathbf{x} | \mathbf{u}, c)$. This can be achieved by a generative model described by the mapping

$$\mathbf{x} = f(\mathbf{z}, \mathbf{u}, c), \quad (3)$$

where $\mathbf{z}$ is a random latent vector following a prior distribution $p(\mathbf{z})$, usually uniform or Gaussian. This generating function $f(\mathbf{z}, \mathbf{u}, c)$ is to be trained with data. This formulation allows differentiation among gNBs while capturing similarities in the RSS spatial distribution.

Exploitation of the Model: Once trained, this generative model can be conveniently applied in simulations. UAV trajectories and gNB locations can be stochastically generated based on a deployment model, providing the condition vector $\mathbf{u}$ for each UAV-gNB link. Random vectors $\mathbf{z}$ can be sampled from the prior distribution, and with $\mathbf{u}$, $\mathbf{z}$, and the categorical information c , the sequence of RSS values $\mathbf{x}$ can be obtained. These RSS values can be generated for intended and interfering links, enabling the computation of many quantities




Fig. 2. Proposed generative channel model architecture.

of interest as experienced by a UAV along its route, say signal-to-interference-plus-noise ratios, bit rates, and frequency and success of handovers (cell reselections).

III. METHODOLOGY

A. Generative Model Architecture

The proposed architecture incorporates two types of conditioning data, namely the distance sequences $\mathbf{u}$ and the gNB indices c . This enables the generation of RSS sequences along specific trajectories for a particular gNB. The architecture, depicted in Fig. 2, is based on the transformer time-series conditional GAN (TTS-CGAN) [17] and multivariate time-series conditional GAN (MTS-CGAN) [18] with appropriate modifications. Specifically, the conditioning on gNB indices follows the auxiliary classifier GAN paradigm of TTS-CGANs while that on distance sequences relies on the classical conditional GAN paradigm of MTS-CGANs. Both the generator (G) and the discriminator (D) share a similar structure. They consist of three TransformerEncoder layers, each with five attention heads. Each layer includes a multi-head attention module followed by a feed-forward multi-layer perceptron (MLP) with a Gaussian error linear unit (GELU) activation function. Normalization layers precede both blocks, and dropout layers (with drop rate $\rho = 0.5$) follow them. Generator and discriminator only differ in their input and output layers. Their structure is summarized in Table I.

Generator: The input consists of a label embedding for the conditioning variable c and a linear transformation of the distance sequence $\mathbf{u}$, to allow for their concatenation together with the random latent vector $\mathbf{z}$. The concatenated vector is then mapped to a sequence of the same length with 50 embedding dimensions. With this configuration, the task is to generate RSS sequences related to a specific gNB, based on the distance sequences experienced from that gNB.

Discriminator: The input module includes a patch and a positional embedding layer. Letting C be the number of gNBs, a sequence of N steps, whether real or generated, can be viewed as an image of shape $(C, 1, W)$, i.e., with height 1,

TABLE I
GENERATOR, DISCRIMINATOR, AND TRAINING PARAMETERS

Generator layers	Number of weights
Distance encoder (Linear + LeakyReLU)	12850
Embedding	386560
TransformerEncoder $\times 3$	1330×3
Channel reduction (Conv2D)	11
Discriminator layers	Number of weights
Patch Embedding	2550
TransformerEncoder $\times 3$	30650×3
Adversarial classification head	151
Training parameter	Value
Generator and discriminator learning rates	$3 \cdot 10^{-4}$ and 10^{-4}
β_1 and β_2	0.9 and 0.999
Batch size and patch size	64 and 16
Number of iterations	$5 \cdot 10^5$

width $W = N$, and C color channels. This image is evenly divided into W/P patches, each with shape $(C, 1, P)$, and a learned positional encoding value is added to each patch to preserve its positional information.

The output module consists of a binary classification layer to classify signals as true or generated, and a multi-class classification layer to determine the originating gNB. Based on the distance sequences provided as conditioning data, the discriminator has two objectives:

- Adversarial classification, correctly classifying time series as true or generated.
- Categorical classification, accurately assigning the label to the input series.

B. Loss Function

The game between discriminator and generator spans two levels: adversarial classification and categorical classification. Let $\hat{\mathbf{x}} = G(\mathbf{z} | \mathbf{u}, c)$ be the generator's output while $D_{\text{adv}}(\cdot)$ is the output of the discriminator's adversarial head and $D_{\text{cls}}(\cdot)$ the output of the discriminator's classification head.

Adversarial Classification: The least-squares (LS) loss function is adopted [19], with the discriminator D minimizing

$$L_{\text{LS}}^{(D)} = \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_X(\mathbf{x})} \left[(D_{\text{adv}}(\mathbf{x} | \mathbf{u}) - t_{\text{R}})^2 \right] + \frac{1}{2} \mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} \left[(D_{\text{adv}}(\hat{\mathbf{x}} | \mathbf{u}) - t_{\text{F}})^2 \right], \quad (4)$$

where t_{R} and t_{F} are the binary labels (targets) for real and generated data, respectively. The generator G minimizes

$$L_{\text{LS}}^{(G)} = \frac{1}{2} \mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} \left[(D_{\text{adv}}(\hat{\mathbf{x}} | \mathbf{u}) - t_{\text{R}})^2 \right], \quad (5)$$

thus attempting to make D classify generated data $\hat{\mathbf{x}}$ as real.

Categorical Classification: Based on the cross-entropy (CE) loss, discriminator and generator respectively minimize

$$L_{\text{CE}}^{(D)} = -\mathbb{E}[\log D_{\text{cls}}(\mathbf{x} | \mathbf{u})] \quad (6)$$

$$L_{\text{CE}}^{(G)} = -\mathbb{E}[\log D_{\text{cls}}(\hat{\mathbf{x}} | \mathbf{u})]. \quad (7)$$

Overall Loss: Discriminator and generator respectively seek the minimum of the total loss functions

$$L_{\text{D}} = L_{\text{LS}}^{(D)} + L_{\text{CE}}^{(D)} \quad (8)$$

$$L_{\text{G}} = L_{\text{LS}}^{(G)} + L_{\text{CE}}^{(G)}. \quad (9)$$

C. Validation Through First- and Second-Order Statistics

To evaluate the model's ability to capture the underlying probability distribution, the first- and second-order statistics of real and generated signals are compared. To that end, the marginal cumulative distribution function (CDF) of the RSS and the correlation matrix distance (CMD) between real and generated RSS sequences are examined.

Marginal Distribution: For a given gNB, consider two sets containing the real and generated RSS sequences, both sets having shape (B, W) with B and W being the batch size and sequence length, respectively. Each set is flattened into a 1D vector of length $n = B \cdot W$ over which the CDF is computed.

Correlation Matrix Distance: The sequence length W is regarded as the number of random variables while B is treated as the number of realizations. Covariance matrices Σ_1 and Σ_2 are then computed for the sets of real and generated sequences, and from those the correlation matrices: from Σ_1 , we construct $\mathbf{A}_1 = \sqrt{\text{diag}(\Sigma_1)}$ and then normalize Σ_1 into

$$\mathbf{R}_1 = \mathbf{A}_1^{-1} \Sigma_1 \mathbf{A}_1^{-1}. \quad (10)$$

Similarly, from Σ_2 we construct $\mathbf{A}_2$ and obtain $\mathbf{R}_2$. Matrices $\mathbf{R}_1$ and $\mathbf{R}_2$ contain every second-order statistic for the real and generated RSS sequences, respectively, with no assumptions on stationarity. (If wide-sense stationarity holds, the matrices are Toeplitz, but in general that is not the case for the RSS.) The CMD between $\mathbf{R}_1$ and $\mathbf{R}_2$ is [20]

$$\text{CMD}(\mathbf{R}_1, \mathbf{R}_2) = 1 - \frac{\text{trace}(\mathbf{R}_1 \mathbf{R}_2)}{\|\mathbf{R}_1\|_{\text{F}} \|\mathbf{R}_2\|_{\text{F}}} \in [0, 1], \quad (11)$$

where $\|\cdot\|_{\text{F}}$ denotes Frobenius norm. The CMD tends to 1 if $\mathbf{R}_1$ and $\mathbf{R}_2$ are maximally different, while it vanishes if they are equal up to a scaling factor, which is the desired outcome.

IV. NUMERICAL RESULTS

Next, the dataset production process is detailed and two case studies are presented to showcase the model's effectiveness.

A. Dataset Production

Due to the limited availability of data on UAV channels, the ray tracing package Wireless InSite by Remcom is employed and the RSS is computed by adding the ray powers at any given location [21, Sec. 4.2].

Deployment Scenario: A 3D representation of a region measuring $1200 \text{ m} \times 1120 \text{ m}$ is imported, corresponding to the city of Boston (see Fig. 1). The representation includes terrain and building data. Transmitting gNBs are manually positioned on three rooftops, 30 m above street level. These sites are potential locations for providing connectivity to UAVs or other aerial devices [7], [22].

Ray Tracing: Simulations are conducted at 28 GHz, which is the dominant frequency for emerging 5G mmWave systems. Buildings are modeled as made of concrete with permittivity 5.31 F/m and conductivity of 0.024 S/m, and the maximum number of reflections and diffractions are set to 6 and 1, respectively. For each gNB, an RSS map is generated over the entire region, sampled every 2 m at a height of 30 m,




Fig. 3. CMD computed on the test set vs. number of training iterations for the case of a single gNB, with and without data augmentation.

and assuming 0 dBm transmit power and unitary antenna gains at both transmitter and receiver. The 2-m sample spacing is a conservative choice that ensures minimal change in the large-scale channel behavior across consecutive samples. (Further modeling the small-scale fading would require a smaller spacing, on the order of half the wavelength, as that is the minimum coherence distance of the small-scale process.)

RSS Sequence Production: UAVs are placed along multiple random trajectories at 30 m of height, produced via MATLAB's Vehicle Network Toolbox, each consisting of 600–800 steps with a 2-m interval. The RSS sequences $\mathbf{X}$ are created by recording the RSS values at each intersection between a trajectory and the power map associated with each gNB. The 3D distance between each intersection and the gNB's location are computed to construct the distance sequences $\mathbf{U}$. Let $\mathbf{X} = [\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}]$ and $\mathbf{U} = [\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(K)}]$ assemble all RSS sequences and all 3D distances, where K is the total number of each. Further, with c identifying the gNB assigned to each sequence, let $\mathbf{c} = [c_1, \dots, c_K]$. Thus, $\mathbf{x}$, $\mathbf{U}$, and $\mathbf{C}$ encapsulate all of the information gathered during the data production process.

Training Data Augmentation: Augmentation is employed to further expand the training data. This facilitates the model's convergence, particularly in cases where data is scarce (say in the case of a single gNB, as discussed next). While this step may not be necessary for ray tracing data, as more trajectories and more gNB power maps could be produced, it becomes crucial when the model is trained with measurements, which are considerably more time-consuming. Using the Python package `tsaug`, a self-convolution operation is applied to each RSS sequence with a flat kernel window of size 20. The resulting convolved sequences are then appended to the existing dataset.

B. Case Study I: Single gNB

To begin with, let us consider only gNB 2 in Fig. 1. Without the need to differentiate among different gNBs, the categorical classification can be disabled and the training can focus solely on the classical adversarial game between generator and discriminator. This entails updating the weights of the discriminator and generator based only on (4) and (5).

Fig. 3 displays the CMD on the test set as a function of the number of training iterations, with and without training data augmentation via convolution. Data augmentation does help to achieve a smaller correlation matrix distance between real and generated RSS sequences, resulting in a more accurate

TABLE II
FINAL CMD VALUES WITH AND WITHOUT DATA AUGMENTATION

Augmentation	CMD with one gNB	CMD with three gNBs
No	0.1609	0.1025
Yes	0.0661	0.0874

and spatially consistent model. The final value of the CMD on the test set is reported in Table II, providing further evidence of the model's accuracy.

Fig. 4 illustrates the evolution of the correlation matrix for the generated RSS sequences as the training progresses, when data augmentation is employed during training, with the conditioning distance sequences extracted from the test set. Note that the rows/columns of Fig. 4 can be interpreted as the autocorrelation of the RSS sequences. The correlation matrix for the corresponding real RSS sequences is also displayed, demonstrating the convergence.

C. Case Study II: Multiple gNBs

Next, let us evaluate the generator's ability to capture different distributions by reintroducing the categorical classification head of the discriminator. The model is trained with data from the three gNBs in Fig. 1, updating the discriminator and generator according to (8) and (9).

Fig. 5 presents, on the left-hand side, the CDF of real and generated RSS values corresponding to the distances in the test set, for gNB2 with data augmentation. On the right-hand side, it presents the mean (solid and dashed lines) and standard deviation (shaded areas) of the real and generated RSS as a function of the UAV-gNB distance. Log-distance least-squares fits are also shown, with the reference distance set to the central point of the axis (2.25 km) and path loss exponents of 2.38 (real) and 2.44 (generated). The model is seen to successfully learn the distribution of the RSS and its dependence on the distance.

Fig. 6 depicts the CMD on the test set, averaged for the three gNBs, as a function of the number of training iterations, with and without data augmentation. As the CMD decreases, the model learns to generate spatially consistent RSS sequences that are stochastically similar to the real ones. Data augmentation aids in achieving an even smaller CMD between real and generated sequences. The final values of the correlation matrix distance are reported in Table II. Fig. 6 also displays the evolution of the classification loss (7) at the generator for the case of data augmentation.¹ As the total loss nears zero, the model successfully discriminates among gNBs.

V. CONCLUSION

This letter has introduced a GNN architecture for spatially consistent air-to-ground channel modeling. The approach effectively captures the spatial dependencies in RSS sequences (equivalently in large-scale channel gains) from multiple gNBs and can be instrumental for system-level evaluations of various metrics. While the presented case studies relied on ray tracing data, the model can also be trained with field measurements.

Follow-up work could aim—at the expense of a faster spatial sampling rate and increased complexity—at incorporating

¹The discriminator classification loss (6), not shown, follows a similar trend.




Fig. 4. Evolution of the correlation matrix for generated RSS sequences as compared to the one of real RSS sequences.




Fig. 5. Left-hand side: CDF of the real and generated RSS values. Right-hand side: mean (solid and dashed lines) plus/minus standard deviation (shaded areas) of the real and generated RSS vs. distance and respective log-distance least-squares fits. Generated sequences are obtained by training with data augmentation.




Fig. 6. CMD averaged across multiple gNBs vs. training iterations. The categorical classification loss of the generator is also shown.

the small-scale fading, altogether delivering the full multipath response with path gains, delays, and angles of arrival and departure for all propagation paths. These would further extend the applicability of the model to multiantenna communication. Exploring the model's ability to generalize across environments would be another relevant research direction.

ACKNOWLEDGMENT

Remcom contributed by providing the Wireless Insite software.

REFERENCES

- [1] G. Geraci et al., "What will the future of UAV cellular communications be? A flight from 5G to 6G," *IEEE Commun. Surveys Tuts.*, vol. 24, no. 3, pp. 1304–1335, 3rd Quart., 2022.
- [2] Q. Wu et al., "A comprehensive overview on 5G-and-beyond networks with UAVs: From communications to sensing and intelligence," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 10, pp. 2912–2945, Oct. 2021.
- [3] Y. Zeng, I. Güvenc, R. Zhang, G. Geraci, and D. W. Matolak, Eds., *UAV Communications for 5G and Beyond*. Hoboken, NJ, USA: Wiley, 2020.
- [4] S. Karimi-Bidhendi, G. Geraci, and H. Jafarkhani, "Optimizing cellular networks for UAV corridors via quantization theory," 2023, *arXiv:2308.01440*.
- [5] A. Garcia-Rodriguez, G. Geraci, D. Lopez-Perez, L. G. Giordano, M. Ding, and E. Bjornson, "The essential guide to Realizing 5G-connected UAVs with massive MIMO," *IEEE Commun. Mag.*, vol. 57, no. 12, pp. 84–90, Dec. 2019.
- [6] M. M. U. Chowdhury, I. Guvenc, W. Saad, and A. Bhuyan, "Ensuring reliable connectivity to cellular-connected UAVs with uptilted antennas and interference coordination," *ITU J. Future Evol. Technol.*, vol. 2, no. 2, pp. 165–185, 2021.
- [7] G. Geraci, D. López-Pérez, M. Benzaghta, and S. Chatzinotas, "Integrating terrestrial and non-terrestrial networks: 3D opportunities and challenges," *IEEE Commun. Mag.*, vol. 61, no. 4, pp. 42–48, Apr. 2023.
- [8] "Study on enhanced LTE support for aerial vehicles; (Release 15)," 3GPP, Sophia Antipolis, France, Rep. TR-36.777, Dec. 2017.
- [9] W. Khawaja, I. Guvenc, D. W. Matolak, U.-C. Fiebig, and N. Schneckenburger, "A survey of air-to-ground propagation channel modeling for unmanned aerial vehicles," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 3, pp. 2361–2391, 3rd Quart., 2019.
- [10] A. A. Khuwaja, Y. Chen, N. Zhao, M.-S. Alouini, and P. Dobbins, "A survey of channel modeling for UAV communications," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 4, pp. 2804–2821, 4th Quart., 2018.
- [11] J. Huang et al., "A big data enabled channel model for 5G wireless communication systems," *IEEE Trans. Big Data*, vol. 6, no. 2, pp. 211–222, Jun. 2020.
- [12] E. Ostlin, H.-J. Zepernick, and H. Suzuki, "Macrocell path-loss prediction using artificial neural networks," *IEEE Trans. Veh. Technol.*, vol. 59, no. 6, pp. 2735–2747, Jul. 2010.
- [13] E. Dall'Anese, S.-J. Kim, and G. B. Giannakis, "Channel gain map tracking via distributed kriging," *IEEE Trans. Veh. Technol.*, vol. 60, no. 3, pp. 1205–1211, Mar. 2011.
- [14] W. Xia et al., "Generative neural network channel modeling for Millimeter-wave UAV communication," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9417–9431, Nov. 2022.
- [15] Y. Yang, Y. Li, W. Zhang, F. Qin, P. Zhu, and C.-X. Wang, "Generative-adversarial-network-based wireless channel modeling: Challenges and opportunities," *IEEE Commun. Mag.*, vol. 57, no. 3, pp. 22–27, Mar. 2019.
- [16] T. J. O'Shea, T. Roy, and N. West, "Approximating the void: Learning stochastic channel models from observation with variational generative adversarial networks," in *Proc. ICNC*, 2019, pp. 681–686.
- [17] X. Li, A. H. H. Ngu, and V. Metsis, "TTS-CGAN: A transformer time-series conditional GAN for biosignal data augmentation," 2022, *arXiv:2206.13676*.
- [18] A. Madane, M.-D. Dilmli, F. Forest, H. Azzag, M. Lebbah, and J. Lacaille, "Transformer-based conditional generative adversarial network for multivariate time series generation," 2022, *arXiv:2210.02089*.
- [19] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. IEEE ICCV*, 2017, pp. 2794–2802.
- [20] M. Herdin, N. Czink, H. Ozcelik, and E. Bonek, "Correlation matrix distance, a meaningful measure for evaluation of non-stationary MIMO channels," in *Proc. IEEE 61st VTC*, 2005, pp. 136–140.
- [21] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Hoboken, NJ, USA: Prentice Hall, 2001.
- [22] S. Kang et al., "Millimeter-wave UAV coverage in urban environments," in *Proc. IEEE Glob. Commun. Conf. (GLOBECOM)*, 2021, pp. 1–6.