

Generative Neural Network Channel Modeling for Millimeter-Wave UAV Communication

William Xia, *Student Member, IEEE*, Sundeep Rangan, *Fellow, IEEE*,
Marco Mezzavilla, *Senior Member, IEEE*, Angel Lozano, *Fellow, IEEE*,
Giovanni Geraci, *Senior Member, IEEE*, Vasilii Semkin,
and Giuseppe Loianno, *Member, IEEE*

Abstract

The millimeter wave bands are being increasingly considered for wireless communication to unmanned aerial vehicles (UAVs). Critical to this undertaking are statistical channel models that describe the distribution of constituent parameters in scenarios of interest. This paper presents a general modeling methodology based on data-training a generative neural network. The proposed generative model has a two-stage structure that first predicts the link state (line-of-sight, non-line-of-sight, or outage), and subsequently feeds this state into a conditional variational autoencoder (VAE) that generates the path losses, delays, and angles of arrival and departure for all the propagation paths. The methodology is demonstrated for 28 GHz air-to-ground channels between UAVs and a cellular system in representative urban environments, with training datasets produced through ray tracing. The demonstration extends to both standard base stations (installed at street level and downtilted) as well as dedicated base stations (mounted on rooftops and uptilted). The proposed approach is able to capture complex statistical relations in the data and it significantly outperforms standard 3GPP models, even after refitting the parameters of those models to the data.

W. Xia, S. Rangan, M. Mezzavilla, and G. Loianno are with NYU Tandon School of Engineering, Brooklyn, USA. Their work is supported by NSF grants 1302336, 1564142, 1547332, and 1824434, NIST, SRC, and the industrial affiliates of NYU WIRELESS.

A. Lozano and G. Geraci are with Univ. Pompeu Fabra, Barcelona. Their work is supported by ERC grant 694974, by MINECO's Project RTI2018-101040, by ICREA, and by the Junior Leader Fellowship Program from "la Caixa" Banking Foundation.

V. Semkin is with VTT Technical Research Centre of Finland Ltd, Finland. His work is supported in part by the Academy of Finland.

Some of the material in this paper was presented at the 2020 IEEE GLOBECOM conference [1].

Index Terms

UAV, drone, mmWave communication, 5G, cellular network, air to ground, channel model, ray tracing, variational autoencoder, generative neural network, 3GPP.

I. INTRODUCTION

Communication with unmanned aerial vehicles (UAVs) is a subject of growing interest, and the millimeter wave (mmWave) range is an inviting realm for this purpose because of the conjunction of enormous bandwidths and line-of-sight (LOS) situations [2]–[17]. As with all communication systems, the design and evaluation of mmWave UAV networks hinges crucially on the availability of suitable channel models.

As current 3rd Generation Partnership Project (3GPP) channel models, which extend up to 100 GHz for terrestrial users, are only calibrated for UAVs at sub-6 GHz frequencies [18], there is a pressing need to extend the availability of channel models suitable for UAVs to the mmWave range. For example, [19] proposes a propagation model for UAV-to-UAV communication at 60 GHz in LOS conditions, and with UAV altitudes ranging between 6 and 15 m. Several other works have also attempted to model various aggregate statistics of the channel model, such as the omnidirectional path loss or narrowband fading [20]–[24]. More generally, as mmWave systems rely on highly directional communication over wide bandwidths, statistical descriptions of the full double directional characteristics of the channel are required, meaning a description of the totality of path components (angles of arrival and departure, gains and delays).

Statistical channel models enable producing random instances of the full set of channel parameters. The joint statistical distribution of these parameters (path angles, gains, and delays) must first be distilled from a combination of physical considerations and field measurements [25], [26], a process that has become increasingly cumbersome as the systems being modelled have grown in complexity and heterogeneity (new frequency bands, broader bandwidths, massive antenna arrays, diverse deployments) [27]. In aerial settings, this complexity is further compounded by additional parameter dependencies on the UAV altitudes, their 3D orientation, or the building heights, among others [18], [28]–[32]. Altogether, the model parameters are bound to exhibit decidedly complex relationships that are difficult to establish through analytical or physical considerations.

Modern data-driven machine-learning methods become an attractive recourse whenever physically based modeling is difficult. Importantly, these methods entail minimal assumptions and

can naturally capture intricate probabilistic relationships. In such spirit, this paper considers data-driven methods to model mmWave air-to-ground channels.

Neural networks (NNs) have been advocated in [33]–[37] for indoor mmWave channel modeling, whereby, upon an input corresponding to some location, the NN outputs the model parameters for that location; in essence, the parameters are then a regression from the training dataset, much as in data-based signal power maps and in learning-based planning and prediction tools [38]–[45]. A strong aspect of all these works is their inherent site-specific nature, a virtue when it comes to optimizing specific deployments. Alternatively, there is interest in models that can produce channel parameters broadly representative of some general environment, say an urban microcellular system.

Generative NNs, which have proven enormously successful with images and text [46]–[48], offer a natural approach to data-driven channel modeling that can broadly represent complex settings, and some early works have successfully trialed generative adversarial networks (GANs) for simple wireless channels [49]–[51]. The present paper propounds a different generative NN structure, powerful and widely applicable, for air-to-ground channel modeling. For data provisioning, we rely on the ray tracing tool [52], which has developed substantially for mmWave communication [53]–[58] and can supply datasets of the size required to train large NNs.

Ray tracing requires a detailed blueprint of the environment, including the size, shape, and location of all obstacles, along with their electromagnetic properties. As it employs high-frequency approximations, ray tracing exhibits some inaccuracies, but is perfectly adequate for our purpose here, which is to validate the proposed methodology. We hasten to emphasize that, ultimately, the model is meant to be driven by field data, gathered either through targeted measurement campaigns or directly supplied by users of the service.

The highlights of this work are as follows:

- *Double-directional wideband characterization.* As chief point, we demonstrate that the proposed method can capture the directional characteristics of the channel at both transmitter and receiver along with its wideband nature, meaning the angular, gain, and delay information for all the paths on each link. This description can be integrated into a standard 3GPP evaluation methodologies [18], [28] and can provide the full wideband MIMO response given specific antenna configurations at transmitter and receiver. No prior assumptions are made regarding the dependencies among parameters, and the model is able to capture relationships that are nuanced and interesting.

- *Novel two-stage structure.* The generative model features a novel two-stage structure where a first NN determines whether the link is in a state of LOS, non-line-of-sight (NLOS), or outage, while a second stage employs a conditional VAE to generate the path parameters given that state. Importantly, several pre-processing steps are introduced to map the path parameters to a format compatible with the NN outputs.
- *Application to UAV mmWave settings.* The methodology is demonstrated by characterizing 28 GHz channels connecting UAVs with two distinct classes of ground base stations. For this application, we demonstrate that trained models can be readily applied to predict complex system parameters such as coverage as a function of antenna pattern and base station deployment.
- *Intra- and inter-environment generalization.* The model is separately trained on data from various environments, namely sections of Tokyo, Beijing, London, Moscow, and Boston. Then, these models are tested on new points from the respective datasets as well as on points from the other datasets. This allows testing the ability of the model trained in one environment to describe the behavior in new locations within that environment (intra-environment generalization) and in locations in other environments (inter-environment generalization).
- *Benchmarking against 3GPP models.* The proposed generative model is benchmarked against the existing 3GPP channel models [18], [28], recalibrated to fit the mmWave data used to train our generative model. In this head-to-head comparison, the generative model proves superior, highlighting the advantage of techniques that make minimal prior structural assumptions.
- *Ability to capture complex statistical relationships.* The comparison to 3GPP models shows that neural networks can capture complex relationships between paths, not described in standard parametric models. For example, we show the model can capture the variation of angular spread and number of paths as a function of the distance.
- *Publicly available model.* The developed model is publicly available [59] and can be readily incorporated to any simulator of mmWave UAV communication that can accept multipath parameters as channel descriptions. And, beyond this use case, the underlying modeling framework may be enticing for other emerging communication scenarios such as terahertz systems, and as an alternative to channel models in more established contexts.

The paper is organized as follows. Section II frames the problem, Section III sets forth the proposed generative approach, and Section IV describes the data procurement process. Then, Section V presents a battery of results that illustrate how the trained model successfully predicts

the channel's behavior in unseen locations. Finally, Section VI contrasts the predictive power of the proposed model against that of the refitted 3GPP model, and Section VII concludes the paper.

II. PROBLEM FORMULATION

We consider the modeling of channels linking a transmitter with a receiver. The UAV is taken to be the transmitter while the base station—gNB in 3GPP terminology [18]—is the receiver, yet, owing to reciprocity, the roles of transmitter and receiver are interchangeable. Each link is described by the collection of parameters [60]

$$\mathbf{x} = \left\{ \left(L_k, \phi_k^{\text{rx}}, \theta_k^{\text{rx}}, \phi_k^{\text{tx}}, \theta_k^{\text{tx}}, \tau_k \right), k = 1, \dots, K \right\}, \quad (1)$$

where K is the number of paths whereas L_k is the loss of path k , $(\phi_k^{\text{rx}}, \theta_k^{\text{rx}})$ are its azimuth and elevation angles of arrival, $(\phi_k^{\text{tx}}, \theta_k^{\text{tx}})$ are its azimuth and elevation angles of departure, and τ_k is its absolute propagation delay. Unlike 3GPP spatial cluster models (e.g., [28]), we do not consider angular or delay dispersion within each path. This is not a limitation of the model, but only a consequence of the tool that produces training datasets with discrete paths; if angular or delay spread information were available, these aspects could be incorporated.

For the sake of specificity, the number of paths is fixed at $K = 20$ with $L_k = L_{\max}$ for paths that are not actually present; we set $L_{\max} = 200$ dB, which is compatible with the maximum path loss detectable by the ray tracer. With these settings, the data vector $\mathbf{x}$ in (1) contains $6K = 120$ parameters per link. Let

$$\mathbf{u} = [\mathbf{d}, c] \quad (2)$$

denote the *link condition*, with $\mathbf{d} = [d_x, d_y, d_z]$ the vector connecting the UAV with the gNB and with c indicating the type of gNB. For air-to-ground modeling, we consider two types of gNBs:

- *Standard gNBs*, installed at street level and downtilted to serve terrestrial users, but potentially usable for UAV connectivity; and
- *Dedicated gNBs*, mounted on rooftops and uptilted, intended specifically for UAVs.

One could also consider other aspects, such as the gNB height, within c . Our methodology is general.

The goal is to capture the conditional distribution $p(\mathbf{x}|\mathbf{u})$, that is, to model the distribution of the paths in a link as a function of that link's condition in some environment. As anticipated, we consider a generative scheme in which

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

where $\mathbf{z}$ is a random vector, termed the *latent vector*, with some fixed prior distribution $p(\mathbf{z})$, while $g(\mathbf{u}, \mathbf{z})$ is the *generating function*, to be trained from data.

Once trained, generative models are conveniently applicable in simulations: the locations of UAVs and gNBs are determined, either deterministically or stochastically according to some deployment strategy, providing the condition vector $\mathbf{u}$ for each link. Random vectors $\mathbf{z}$ can then be produced for each link from the prior $p(\mathbf{z})$ and, from $\mathbf{u}$ and $\mathbf{z}$, the path parameters $\mathbf{x}$ follow as per (3). These parameters can be generated for both intended and interfering links and, in conjunction with the antenna patterns, array configuration, and beam tracking methods, allow computing quantities of interest such as signal-to-noise ratios (SNRs), signal-to-interference-plus-noise ratios (SINRs), or bit rates.

Small-scale dynamics can also be modeled under the premise of local stationarity. Specifically, given any local motion with some velocity, Doppler shifts can be computed and applied to each path to derive the time-varying wideband frequency response [60]. Statistical modeling of large-scale dynamics such as blockage [61], [62] and spatial consistency [28] remain an interesting avenue of future research.

III. PROPOSED GENERATIVE MODEL

A. Overview

The propounded generative model, sketched in Fig. 1, constructs the generative function as two cascaded stages, namely a link-state prediction stage followed by a path generation stage. The latent vector $\mathbf{z}$ subsumes three components,

$$\mathbf{z} = [z_{\text{state}}, \mathbf{z}_{\text{NLOS}}, \mathbf{z}_{\text{out}}]. \quad (4)$$

The link-state predictor accepts the condition vector $\mathbf{u}$ and a random variable z_{state} , from which it determines the link state s . From s and the two other latent components, $\mathbf{z}_{\text{NLOS}}$ and $\mathbf{z}_{\text{out}}$, the path generation stage then produces the final path parameters $\mathbf{x}$. We next describe the details of this whole architecture.




Fig. 1: Overall architecture for the two-stage generative model, which accepts a link condition vector $\mathbf{u}$ and a latent vector $\mathbf{z} = [z_{\text{state}}, \mathbf{z}_{\text{NLOS}}, \mathbf{z}_{\text{out}}]$ to generate random path parameters $\mathbf{x}$. (For the sake of clarity, various transformations, described in the text, are omitted from this diagram.)

B. Link-State Predictor

As recognized by 3GPP models such as [28], it is crucial to first determine the existence or lack of the LOS path. To this end, the link-state predictor accepts the condition $\mathbf{u}$ defined in (2) and produces probabilities for the link being in one of three states [63]:

- LOS: The LOS path is present, possibly in addition to NLOS paths;
- NLOS: The LOS path is blocked, but at least one NLOS path is active;
- NoLink: No propagation paths (either LOS or NLOS) exist for this link.

In the sequel, $s \in \{\text{LOS}, \text{NLOS}, \text{NoLink}\}$ denotes the predicted link state while the generative model mapping $\mathbf{u}$ to s is represented by

$$s = g_{\text{state}}(\mathbf{u}, z_{\text{state}}). \quad (5)$$

Such mapping entails three steps, expounded next.

1) *Condition Vector Transformation*: The vector $\mathbf{u}$ is transformed into a new vector

$$\left[\mathbf{c}_{\text{one}}, d_{3\text{D}} \mathbb{1}_{\{c=1\}}, d_z \mathbb{1}_{\{c=1\}}, \dots, d_{3\text{D}} \mathbb{1}_{\{c=C\}}, d_z \mathbb{1}_{\{c=C\}} \right] \quad (6)$$

where $\mathbf{c}_{\text{one}}$ is a one-hot coded version of the gNB type c while d_z is the vertical distance,

$$d_{3\text{D}} = \sqrt{d_x^2 + d_y^2 + d_z^2} \quad (7)$$

is the 3D distance, C is the number of possible gNB types, and $\mathbb{1}_{\{c=i\}}$ is the indicator function for the event $c = i$. As c can take C possible values, we can one-hot code $\mathbf{c}_{\text{one}}$ with $C - 1$ dimensions. Hence, the transformed vector in (6) has dimension $C - 1 + 2C = 3C - 1$. With $C = 2$ (standard and dedicated), the transformed vector in (6) has $3C - 1 = 5$ components. The motivation for the transformation in (6) is to enable a different behavior of the first layer of the NN for different types of gNB.

The transformed vector in (6) is passed through a min-max scaler that maps its components to values between 0 and 1; the limits on this min-max scaler are learned during training. The resulting transformed and scaled value is denoted by $\mathbf{v}_{\text{state}}$.

2) *NN*: A fully connected NN, configured as per Table I, generates the link-state probabilities. The input to this NN is $\mathbf{v}_{\text{state}}$ while its output is a three-way softmax corresponding to the three states.

3) *Sampling*: In the final step, a uniform random variable $z_{\text{state}} \in [0, 1]$ samples the link state s based on the probability outputs from the NN.

TABLE I: Generative model configuration

	Link state prediction	Path VAE encoder	Path VAE decoder
Number of inputs	5	5 + 120	5 + 20
Hidden units	[25, 10]	[200, 80]	[80, 200]
Number of outputs	3	20 + 20	120 + 120
Optimizer	Adam	Adam	
Learning rate	10^{-3}	10^{-4}	
Epochs	50	10000	
Batch size	100	100	
Number of NN parameters	1653	44520	40720

C. Path Generation Stage

The second stage generates the parameters $\mathbf{x}$ in (1) given $\mathbf{u}$ and s . This also entails various steps, described next.

1) *Condition Vector Transformation*: Again, we begin by transforming $\mathbf{u}$ and s , in this case into

$$\left[\mathbf{c}_{\text{one}}, d_{3D}, 10 \log_{10}(d_{3D}), d_z, s \right], \quad (8)$$

where $\mathbf{c}_{\text{one}}$, d_{3D} and d_z are as in (6). For this condition vector, we found that including both d_{3D} and $\log_{10}(d_{3d})$ enabled better modeling with a smaller NN. This five-dimensional vector is then passed through a min-max scaler to produce a five-dimensional vector with values between 0 and 1. We denote this transformed vector by $\mathbf{v}_{\text{path}}$.

2) *NLOS VAE*: The next, and most intricate step, is to generate the parameters for the NLOS paths within $\mathbf{x}$. As explained below, these NLOS paths are represented in a transformed version denoted by $\mathbf{y}_{\text{NLOS}}$. For now, we recall that there are up to $K = 20$ NLOS paths with 6 parameters per path, meaning that $\mathbf{y}_{\text{NLOS}}$ is of dimension $6K = 120$.

We want to generate $\mathbf{y}_{\text{NLOS}}$ from $\mathbf{v}_{\text{path}}$ and from some randomness. This mapping should be trained such that the conditional distribution of $\mathbf{y}_{\text{NLOS}}$ given $\mathbf{v}_{\text{path}}$ matches the distribution in the training dataset. There are a large number of methods for training generative models, the two most common being variants of GANs [46], [47] or VAEs [48]. We found the most success with a VAE, as it avoids the minimax optimization required by a GAN.

We apply a standard VAE architecture [48] that has itself two stages: the first stage accepts as inputs a random vector $\mathbf{z}_{\text{NLOS}}$ along with $\mathbf{v}_{\text{path}}$ and it outputs means and variances for the NLOS components, namely

$$[\boldsymbol{\mu}_y, \boldsymbol{\sigma}_y^2] = g_{\text{NLOS}}(\mathbf{v}_{\text{path}}, \mathbf{z}_{\text{NLOS}}). \quad (9)$$

The vectors $\boldsymbol{\mu}_y$ and $\boldsymbol{\sigma}_y^2$ share the dimensions of the sought $\mathbf{y}_{\text{NLOS}}$, hence they combine into 120+120 output values. The entries of $\mathbf{z}_{\text{NLOS}}$ are i.i.d. Gaussian with mean zero and unit variance. In VAE terminology, the dimension of $\mathbf{z}_{\text{NLOS}}$ is termed the *latent dimension*, with higher such dimensions enabling better fitting to the data but requiring larger training datasets. In the remainder, the latent dimension is kept at 20.

The sought $\mathbf{y}_{\text{NLOS}}$ is sampled from the means and variances,

$$\mathbf{y}_{\text{NLOS}} = \boldsymbol{\mu}_y + \boldsymbol{\sigma}_y \odot \mathbf{z}_{\text{out}}, \quad (10)$$

where $\mathbf{z}_{\text{out}}$ has 120 zero-mean unit-variance i.i.d. Gaussian entries and $\odot$ indicates entry-wise multiplication.

In the VAE paradigm, the generator in (9) is termed the *decoder*. The VAE also requires training a so-called *encoder* that maps data samples $\mathbf{y}_{\text{NLOS}}$ and $\mathbf{v}_{\text{path}}$ back to the latent vector $\mathbf{z}_{\text{NLOS}}$. This encoder attempts to approximate sampling from the posterior density of $\mathbf{z}_{\text{NLOS}}$ given $\mathbf{y}_{\text{NLOS}}$ and $\mathbf{v}_{\text{path}}$. The encoder and decoder are then jointly optimized to maximize an approximation of the log-likelihood called the evidence lower bound (ELBO); see [48] for details.

Similar to standard VAE architectures [48], we approximate the posterior density of $\mathbf{z}_{\text{NLOS}}$ given $\mathbf{y}_{\text{NLOS}}$ and $\mathbf{v}_{\text{path}}$ by a Gaussian with independent components. Hence, the encoder takes as inputs $\mathbf{y}_{\text{NLOS}}$ and $\mathbf{v}_{\text{path}}$, and output a vector of mean and a vector of variances for the latent variables $\mathbf{z}_{\text{NLOS}}$. Under this assumption, the encoder can be represented as a function

$$[\boldsymbol{\mu}_z, \boldsymbol{\sigma}_z^2] = h_{\text{NLOS}}(\mathbf{v}_{\text{path}}, \mathbf{y}_{\text{NLOS}}), \quad (11)$$

that takes as inputs $\mathbf{y}_{\text{NLOS}}$ and $\mathbf{v}_{\text{path}}$ and outputs vectors $\boldsymbol{\mu}_z$ and $\boldsymbol{\sigma}_z^2$ representing the mean and variance of $\mathbf{z}_{\text{NLOS}}$ given $\mathbf{y}_{\text{NLOS}}$ and $\mathbf{v}_{\text{path}}$. The vectors $\boldsymbol{\mu}_z$ and $\boldsymbol{\sigma}_z^2$ will have the same dimension as the latent vector $\mathbf{z}_{\text{NLOS}}$. Given the outputs of the encoder, we can then sample from the approximate posterior density by

$$\mathbf{z}_{\text{NLOS}} = \boldsymbol{\mu}_z + \boldsymbol{\sigma}_z \odot \boldsymbol{\epsilon}, \quad (12)$$

where, again, $\odot$ represents elementwise multiplication and $\boldsymbol{\epsilon}$ is i.i.d. zero-mean unit-variance Gaussian noise.

In our case, the encoder and decoder are embodied by fully connected NNs configured as per Table I. Since the latent vector $\mathbf{z}_{\text{NLOS}}$ is realized as a 20-dimensional Gaussian vector, the decoder accepts this 20-dimensional Gaussian vector plus the five-dimensional vector $\mathbf{v}_{\text{path}}$ and yields the 120+120 means and variances. Conversely, the encoder is fed $\mathbf{v}_{\text{path}}$ and a 120-dimensional data input and produces means and variances for the 20-dimensional latent vector.

3) *NLOS Transformation*: As advanced, the generated vector $\mathbf{y}_{\text{NLOS}}$ is a transformed version of the path parameters, the reason being that those actual parameters are heterogeneous: they include path losses, angles, and delays. To put them on an equal footing, $\mathbf{x}_{\text{NLOS}}$ maps onto $\mathbf{y}_{\text{NLOS}}$ as follows:

- The path losses are converted to dB-scale path gains and the minimum such value in the dataset is subtracted out. The resulting excess path gains are then run through a min-max scaler to lie between 0 and 1; a value of zero corresponds to the maximum path loss (L_{max}) and hence to absence of this path altogether.
- The angles are rotated relative to the LOS direction, and then scaled such that 180° corresponds to a unit value.
- The LOS delay is subtracted from the rest of delays, and the resulting excess delays are again scaled to be between 0 and 1.

The above transformations ensure that all values are in a similar range and referenced to the LOS path. The min-max scalers for the path losses and delays are fit to the training data, and

TABLE II: Environment comparisons

Building height	Beijing, China	London, UK	Boston, USA
Minimum	3 m	3 m	2 m
Maximum	238 m	283 m	190 m
Mean	99.9 m	73.6 m	42.4 m
Median	85.5 m	62.5 m	32.8 m

we note that the mapping of angles and delays relative to the LOS path can take place even if such LOS path does not exist (because of blockage).

Once $\mathbf{y}_{\text{NLOS}}$ has been generated, the transformation must be undone to obtain the NLOS path parameters, $\mathbf{x}_{\text{NLOS}}$.

4) *Addition of the LOS Path:* For the LOS path, when it exists, the delay and angles of departure and arrival can be computed from sheer geometry while its loss can be computed from Friis' law [60]. The final step is the addition, when it exists, of such LOS path to $\mathbf{x}_{\text{NLOS}}$, which renders the full collection of path parameters, $\mathbf{x}$.

IV. RAY TRACING DATA AT 28 GHZ

Experimental data on UAV channels is limited, particularly in the mmWave bands [29], [30], [64]–[66]. In this work, we employ a powerful ray tracing package, Wireless InSite by Remcom [52], also used in [53], [57]. For our data production, we consider sections of five cities (Tokyo, Beijing, London, Moscow, and Boston) having varying sizes and distinct types of terrain, buildings, and foliage. Shown in Fig. 2 are 3D representations of these city sections, whose blueprints are part of the Wireless Insite package.

It is standard practice to differentiate channel models across environments, e.g., the 3GPP mmWave model provides separate parameter distributions for distinct types of environment, say *urban macro* and *urban micro* [28]. In a data-driven approach, environment-specific models can be created by appropriately partitioning the training data. In our case, it is natural to define a distinct environment for each of the five represented city sections as these have been selected because of their different characteristics in the first place. This is emphasized by the statistics provided in Table II. We see, for instance, that the section of Boston has a significantly lower average building height than the others and a lower maximum building height as well. We can use this to infer how each city may perform relative to the rest. Furthermore, we can observe




Fig. 2: 3D representations of the five considered city sections: (a) Tokyo, (b) Beijing, (c) London, (d) Moscow, and (e) Boston.

other diversifying aspects of each city by means of a visual inspection. For example, Moscow has buildings that are all nearly the same height and considerable open areas. Although Tokyo has a comparable amount of open areas, those are concentrated in a single quadrant, in contrast with the scattered nature of Moscow's open areas. Because of this, we can expect Moscow to support better coverage in terms of LOS probability and path loss.

The number of deployed transmitters (UAVs) and receivers (gNBs) is detailed in Table III for each of the environments. As advanced in earlier sections, two distinct types of gNBs are manually placed:

- *Standard gNBs*. These are placed on streets at a height of 2 m, emulating typical locations for

TABLE III: City sections and deployment parameters

	Tokyo, Japan	Beijing, China	London, UK	Moscow, Russia	Boston, USA
Area (m^2)	1420×1440	1650×1440	1500×1480	1440×1380	1130×1220
Number of UAVs	140	120	120	160	138
Number of standard gNBs	220	180	122	200	95
Number of dedicated gNBs	200	120	93	160	78

5G microcells intended to serve ground users.

- *Dedicated gNBs*. These are located on rooftops, 30 m above street level, meant to provide additional coverage to UAVs.

Transmitting UAVs, for their part, are placed at different horizontal locations in each environment, at one of four possible altitudes: 30, 60, 90 and 120 m.

In total, 58800 UAV-gNB links are created for the Tokyo environment, 36000 for Beijing, 25800 for London, 57600 for Moscow, and 23874 for Boston. The Wireless InSite ray tracing tool is then run to simulate the channel on every link, producing the path parameters $\mathbf{x}$ for each link (directions of arrival and departure, path losses, and delays). All simulations are conducted at 28 GHz. The maximum number of reflections is set to six per ray and the maximum number of diffractions is set to one. This means that, if a diffraction is present, the ray tracer uses additional logic to determine whether reflections can occur before, after, and between diffractions along the path between Transmitter (TX) and Receiver (RX). The materials of different objects in the 3D model can also be specified, and both ground and wall surfaces are taken to be made of concrete with a permittivity of 5.31 F/m.

The datasets thus gathered are utilized to train the model described in Sec. III.

V. MODELING RESULTS

This section describes various features of the learned models, and their ability to capture interesting wireless phenomena. We also seek to evaluate the generalization ability of the models, meaning their ability to accurately describe the channel behavior in locations other than those in the training dataset. As mentioned in the introduction, this ability is a highly desirable attribute, and hence we test it extensively.

The links available for each environment are split, 75% for training and 25% for testing. Models are then trained separately for each environment, which enables assessing the generalization ability in these two senses:

- *Intra-environment.* The model trained on the 75% training links of a specific dataset is evaluated on the 25% test links of that same dataset. This appraises the ability of the model to generalize to links in the same environment, but at new locations not seen during training.
- *Inter-environment.* The model trained on a specific dataset is evaluated on another dataset. This serves to examine the model's ability to generalize to links in other environments.

All the implementations are based on Tensorflow 2.2; the code, datasets, and pre-trained models can be found in [59].

A. LOS Probability

To illustrate the functioning of the link-state predictor, Fig. 3 shows the probability of the link being in the LOS state as a function of the horizontal distance between UAV and gNB. Precisely, Fig. 3a depicts the actual probabilities in the test data for each of the environments and Fig. 3b depicts the respective model predictions. In both cases, the results are averaged over the four possible UAV altitudes. The link-state predictor is seen to accurately determine the trends in the test data for each of the environments and to reflect the very different behaviors of standard and dedicated gNBs. We also observe interesting differences across environments. The LOS probability is uniformly higher in Moscow, both for standard and dedicated gNBs, consistent with the relatively shorter buildings therein. Beijing, in turn, exhibits a relatively high LOS probability for standard gNBs yet a relatively low LOS probability for dedicated gNBs, a contrast that points to an abundance of both reflection opportunities and blockages.

Insights on the impact of the UAV altitude can be drawn from Figs. 4a and 4b, where again we see the excellent match between the test data and the model predictions thereon. Dedicated gNBs can provide substantially higher probabilities of LOS coverage at long horizontal distances provided the UAV is high enough. In contrast, standard gNBs tend to be far more limited in terms of horizontal coverage.

We will see next how all of the above has a significant impact in other features such as the path loss and path angular distributions.




Fig. 3: LOS probability in the test locations, averaged over the four possible UAV altitudes (30, 60, 90, and 120 m); (a) ray-tracing data, (b) generated by the model.

B. Path Loss: Intra-Environment Evaluation

We now turn to evaluating the accuracy of the rest of the parameters. Fundamentally, we want to measure how close the distribution of the trained generative model in (3) is to the observed conditional distribution of the test data itself. To this end, let $(\mathbf{u}_i, \mathbf{x}_i)$, $i = 1, \dots, N_{ts}$ be the test samples, each containing a link condition, $\mathbf{u}_i$, and its corresponding path parameters, $\mathbf{x}_i$. To evaluate how closely the learned model fits this test data, for each sample we can compute




Fig. 4: LOS probability for (a) Moscow, Russia and (b) Tokyo, Japan, parameterized by altitude and horizontal distance.




Fig. 5: CDF of the path loss for (a) Boston and (b) Moscow.

some statistic $\phi(\mathbf{u}_i, \mathbf{x}_i)$ that is of relevance. As an example of statistic, we compute the path loss experienced by UAVs and gNBs equipped with omnidirectional antennas, deferring to later in the paper the consideration of directivity. Using the same conditions $\mathbf{u}_i$ in the test data, we generate a sample $\mathbf{x}_i^{\text{rnd}} = g(\mathbf{u}_i, \mathbf{z}_i)$ from the trained generative model and some random $\mathbf{z}_i$. We can then compute $\phi(\mathbf{u}_i, \mathbf{x}_i^{\text{rnd}})$ and compare its CDF with that of the actual $\phi(\mathbf{u}_i, \mathbf{x}_i)$.

We first evaluate the intra-environment accuracy of the omnidirectional path loss predictions. Fig. 5 shows the CDF of path losses for the test data of a couple of environments alongside the CDF of path losses generated by the trained model using the same condition values as the

test data. An excellent match is observed for both standard and dedicated gNBs. In particular, the trained generative model is able to capture the multi-slope behavior that arises in some environments due to the mixture of LOS/NLOS links.

C. Path Loss: Inter-Environment Evaluation

Next, we gauge the model's ability to make predictions on an environment after having been trained on a different one. Presented in Fig. 6 are two sets of contrasting such results. On the left-hand side we have the CDF of the path loss on the Moscow test data as predicted by a model trained with the Beijing dataset. For standard gNBs the match is satisfactory, indicating similarity in the respective propagation mechanisms for those gNBs, chiefly reflections. For dedicated gNBs, conversely, the Beijing model largely overstates the Moscow path loss, pointing to important discrepancies in the degree of blockage between the two environments. These observations are fully consistent with those made in Section V-A for the LOS probabilities in Moscow and Beijing. On the right-hand side of the figure, the same exercise is repeated for a model trained with London data and tested in the Tokyo environment, and in this case the agreement is excellent for both standard and dedicated gNBs. We thus see how the proposed methodology enables assessing the inter-environmental generalizability of models, which turns out to depend not only on the environments, but further on the type of gNB. The similarities and




Fig. 6: Inter-environment comparisons for (a) a model that fails to accurately predict the path loss on an environment different than the training one, and (b) a model that does accurately make that prediction.

discrepancies thereby revealed are valuable and highly non-obvious from a visual inspection of the environments in Fig. 2.

Fig. 6a shows a somewhat worse match (CDFs are further apart) for the standard case than what can be observed for its counterpart in Fig. 6b, but is still a far better inter-environment generalization than for the dedicated case. We infer that loss models for standard deployments generalize better to other environments, in comparison with models for dedicated deployments. Besides the two figures that are herein presented supporting this observation, we note that this is a clear pattern throughout the results. The dedicated model might be harder to generalize due to varying building heights and densities, and this is an interesting phenomenon that calls for further research.

D. Angular Distribution

Let us now turn to the path angles. Fig. 7 plots the distribution of those angles as a function of the 3D distance between the UAV and gNB. The distribution is computed over all the paths within 30 dB of the strongest path within each link for all the links in the test dataset, and it is averaged over the four possible UAV altitudes. (For the sake of readability, the links to standard and dedicated gNBs are combined, but respective plots for the standard and dedicated gNBs, or plots to separately observe the effects of elevation and horizontal distance, could just as well be produced.)

Each row in Fig. 7 shows the distribution of one of the angles, $\phi_k^{\text{rx}}, \theta_k^{\text{rx}}, \phi_k^{\text{tx}}, \theta_k^{\text{tx}}$, relative to the LOS direction (even when the LOS path is blocked). For each environment, the left-hand-side column is the distribution for the test data whereas the right-hand-side column is the counterpart generated by the learned model.

The model matches very well the actual angular distribution in the test data. In particular, it captures an important phenomenon: for all distances and angles, the NLOS paths tend to be angularly close to the LOS direction. Moreover, the angular spread decreases as the UAV and gNB are further apart. This behavior makes intuitive sense in that, as the UAV pulls away from the gNB, there is less local scattering to create angular dispersion. Consistent with this, the scattering is much wider at the gNB end of the links.



Fig. 7: Distribution of angles, averaged over the four UAV altitudes, for (a) Tokyo, (b) Beijing, (c) London, (d) Moscow, and (e) Boston.

TABLE IV: Uplink single-cell simulation parameters

Item	Value
Spectrum	Carrier frequency: 28 GHz
	Bandwidth: 400 MHz (4×100 MHz aggregation)
gNB height	Standard: 2 m; Dedicated: 30 m
Array size	UAV: $N_{\text{UAV}} = 16$ (4×4 UPA)
	gNB: $N_{\text{gNB}} = 64$ (8×8 UPA)
Antenna spacing	Half-wavelength
Array vertical orientation	UAV: $180^\circ \downarrow$ lower hemisphere coverage [67]
	Standard gNB: $100^\circ \searrow$ ground coverage, 3 sectors
	Dedicated gNB: $0^\circ \uparrow$ upper hemisphere coverage
Transmit power	UAV: 23 dBm
Losses	6 dB including noise figure [69], [70]

E. SNR Predictions

We finalize by demonstrating a specific application enabled by the generative model. Specifically, we compute the predicted uplink (UAV to gNB) local-average SNR as a function of the UAV position in the single-cell scenario described in Table IV, which is consistent with 5G deployments at 28 GHz [67]. Such uplink SNR is of particular interest since this is usually the power-limited link direction, and the one envisioned to support high-bit-rate applications [5], [6].

A gNB is located at $(0, 0, h)$ with $h = 2$ m and $h = 30$ m in the standard and dedicated cases, respectively. In the standard case, the gNB features three sectors with a half-power beamwidth of 90° per sector and a 100° downtilt (relative to vertical), as customary to serve ground users. Hence, the connections from UAVs to standard gNBs must necessarily be through sidelobes or reflected paths [11], [18]. In the dedicated case, the gNB is single-sectored with an upward-facing array intended for aerial coverage. The UAV, equipped with a single array at its bottom, designed for lower-hemisphere coverage [67], is at $(x, 0, z)$ with $x \in [0, 500]$ m and $z \in [0, 130]$ m. For each UAV position and gNB type (standard or dedicated), 100 channels realizations are generated by the model and used to compute the local-average SNR [68]. The median of these local-average SNRs is depicted in Fig. 8.

The experiment shows how the SNR at any location within the environment can be predicted from the model and the specifics of the setting (arrays, powers, and the other details in Table IV).




Fig. 8: Median predicted local-average SNR as a function of the UAV position for (a) Tokyo, (b) Beijing, (c) London, (d) Moscow, and (e) Boston. The horizontal lines indicate the altitude of the dedicated gNBs.




Fig. 9: Level curves for $\text{SNR} = 10$ dB across all five environments.

The dedicated gNBs provide much better coverage at large horizontal distances, yet standard gNBs can provide solid coverage when the horizontal distance is small (below roughly 100 m). This coverage from standard gNBs is rather surprising: complying with 3GPP specifications [28], standard gNBs have downtilted antennas with a 30 dB front-to-back gain ratio, which hinder the connectivity from direct vertical paths. However, the learned model captures reflections from neighboring buildings within the antenna beamwidth, and the simulations show that these reflected paths do enable coverage.

Another perspective on these results is provided in Fig. 9, which depicts the (2D distance, elevation) contours at which the median SNR equals 10 dB for each environment. The coverage at this SNR is seen to be superior when the UAV is horizontally close to the gNB and/or vertically high (due to increased probability of maintaining LOS). An analysis of Fig. 9 reveals interesting relations between the physical characteristics of each environment and their respective wireless channels. An as example, for a dedicated deployment, we can surmise that Beijing is the best environment at lower altitudes, up to 150 m horizontally; this relates to its sparse building density, as per Fig. 2b. Beyond 500 m horizontally, however, it is Moscow that best supports a dedicated deployment, maintaining a certain SNR at low altitudes; this agrees with our initial predictions about Moscow's coverage based on the height profile of its buildings and its large open areas.

VI. BECHMARKING AGAINST 3GPP MODELS

To complete the test-drive of the proposed generative model, it is fitting to benchmark it against the alternatives offered by existing standards. We focus here on the 3GPP UMi-AV (urban micro with aerial vehicles) scenario [18], [28], which is the closest to our work. This model is suitable for mmWave frequencies up to 100 GHz, but its parameters are calibrated for UAVs—meaning users at altitudes above 22.5 m—only below 6 GHz. In order to provide a fair benchmark for our proposed architecture, we refit those parameters with the data for each of our environments, and restrict the comparisons to standard gNBs.

A. LOS Probability

We first examine the probability of LOS, P_{LOS} . The 3GPP model takes in several parameters such as the heights of transmitter and receiver and, as well as their horizontal distance, $d_{2D} = \sqrt{d_x^2 + d_y^2}$ [28]. If the UAV height h is between 1.5 m and 22.5 m, then

$$P_{\text{LOS}} = \begin{cases} 1 & d_{2D} \leq \alpha_1 \\ \frac{\alpha_1}{d_{2D}} + e^{-\frac{d_{2D}}{\alpha_2}} \left(1 - \frac{\alpha_1}{d_{2D}}\right) & \alpha_1 \leq d_{2D} \end{cases} \quad (13)$$

whereas, if $h > 22.5$ m,

$$P_{\text{LOS}} = \begin{cases} 1 & d_{2D} \leq d_1 \\ \frac{d_1}{d_{2D}} + e^{-\frac{d_{2D}}{p_1}} \left(1 - \frac{d_1}{d_{2D}}\right) & d_{2D} > d_1 \end{cases} \quad (14)$$

with

$$\begin{aligned} p_1 &= \alpha_3 \log_{10}(h) + \alpha_4 \\ d_1 &= \max(\alpha_5 \log_{10}(h) + \alpha_6, \alpha_1). \end{aligned} \quad (15)$$

The values for the parameters, which in the 3GPP model [28] are

$$\begin{aligned} \boldsymbol{\alpha}_{\text{LOS}} &= [\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6] \\ &= [18, 36, 294.05, -432.94, 233.98, -0.95], \end{aligned} \quad (16)$$

are herein refitted for each environment. Specifically, for each link we specify the set of condition variables

$$\mathbf{u} = [\log_{10}(h), d_{2D}, h, h_{\text{gNB}}], \quad (17)$$

and a binary label $y = 1$ if the link is LOS and $y = 0$ otherwise. The 3GPP model can be viewed as a function

$$P_{\text{LOS}} = P(y = 1|\mathbf{u}) \quad (18)$$

$$= g_{\text{LOS}}(\mathbf{u}, \boldsymbol{\alpha}_{\text{LOS}}), \quad (19)$$

mapping the condition vector $\mathbf{u}$ to the LOS probability. From the links $(\mathbf{u}_i, y_i)$ on a given environment, the parameters $\boldsymbol{\alpha}_{\text{LOS}}$ can be found by minimizing the binary cross entropy,

$$J(\boldsymbol{\alpha}_{\text{LOS}}) = - \sum_i \left[y_i \log(g_{\text{LOS}}(\mathbf{u}_i, \boldsymbol{\alpha}_{\text{LOS}})) \right. \quad (20)$$

$$\left. + (1 - y_i) \log(1 - g_{\text{LOS}}(\mathbf{u}_i, \boldsymbol{\alpha}_{\text{LOS}})) \right]. \quad (21)$$

This minimization, which is tantamount to a maximum likelihood estimation of $\boldsymbol{\alpha}_{\text{LOS}}$, is performed via stochastic gradient descent (see Table V). A distinct set of refitted parameters is

TABLE V: 3GPP Refitting Optimization

	P_{LOS}	Path Loss
Number of inputs	5	6
Number of parameters	6	19
Optimizer	Adam	
Loss function	Binary cross entropy	Mean-squared error
Learning rate	10^{-3}	
Epochs	50	
Batch size	128	

obtained for each of the environments in Fig. 2, with the imposition that those parameters are within a multiplicative interval $[0.01, 10]$ of the nominal 3GPP values in (16) to prevent overfitting.

Our proposed generative approach can now be validated against the default 3GPP model and its refitted version. The horizontal distance, d_{2D} , and the vertical distance, d_z , are binned into sections of 20 m and 5 m, respectively. From a histogram of the test links' LOS condition over the bins, the empirical P_{LOS} of the test data is obtained and contrasted with the prediction from the three models.

Table VI shows the mean absolute error of the LOS probability over the grid. We observe that the refitted 3GPP model is significantly better than its default form and that our proposed

TABLE VI: P_{LOS} : Mean Absolute Error

	Tokyo, Japan	Beijing, China	London, UK	Moscow, Russia	Boston, USA
Default 3GPP Model	0.272	0.247	0.303	0.180	0.336
Refitted 3GPP Model	0.040	0.056	0.057	0.058	0.047
Proposed Generative Model	0.036	0.058	0.057	0.034	0.041

approach, with minimal prior structure, performs similarly or better—sometimes markedly—than even the refitted 3GPP model on every environment.

B. Path Loss

We employ a similar strategy to refit the 3GPP path loss model. It is important to note that, unlike with P_{LOS} , which depends exclusively on the geometry, the path loss is frequency-dependent. Separately for the LOS and NLOS cases, the 3GPP model accepts as input the condition vector in (17) and outputs a predicted path loss as some function,

$$\text{PL} = g_{\text{PL}}(\mathbf{x}, \boldsymbol{\alpha}_{\text{PL}}), \quad (22)$$

for specific parameters $\boldsymbol{\alpha}_{\text{PL}}$ whose nominal details are given in [18, Table B-1]. We refit the model, this time with a mean-squared error criterion (see Table V for details).




Fig. 10: Path loss fitting for Tokyo, Japan.

TABLE VII: Path loss: Wasserstein-1 distance to test data distribution (dB)

	Tokyo, Japan	Beijing, China	London, UK	Moscow, Russia	Boston, USA
Default 3GPP Model	15.8	18.8	17.8	15.5	21.4
Refitted 3GPP Model	10.7	14.8	12.3	14.3	14.9
Proposed Generative Model	6.70	2.22	2.95	2.49	3.42

An example is presented in Fig. 10, which depicts the CDFs of path losses for the test data, the default 3GPP model, the refitted 3GPP model, and our proposed approach for Tokyo specifically. While the refitted 3GPP model performs decidedly better than the default one, our proposed approach best approximates the distribution of the actual data. To quantify the differences among the distributions in this and the other environments, we invoke the Wasserstein-1 distance [71]. This metric can be understood as a type of "distance function" between two distributions, and maintains the same units as the random variable (dB). For two distributions P and Q , the Wasserstein distance equals

$$W(P, Q) = \max_f \left[\mathbb{E}(f(X)|X \sim P) - \mathbb{E}(f(X)|X \sim Q) \right], \quad (23)$$

where the maximization is over all Lipschitz functions satisfying $\|\nabla f(x)\| \leq 1 \ \forall x$. This metric is commonly used to train GANs [72] and, for scalar random variables, it can be computed efficiently as the integrated difference in CDFs [73]. Table VII shows the Wasserstein distance between the test data and the various models, with the propounded approach outperforming the rest in every environment.

VII. CONCLUSION

In an age in which the complexity of wireless systems has exploded, generative models are a fitting engine for statistical channel modeling. This is particularly true for the elaborate settings encountered in mmWave UAV communication. Under the proviso that abundant data is available, generative models are perfectly equipped to learn intricate probabilistic relationships and then produce parameters distributed accordingly. The only assumption is the choice of the parameters themselves, which can rest on basic principles of radio propagation.

The proposed generative model, publicly available [59], has been shown to learn effectively and it can hence be calibrated for any desired operating frequency, type of deployment, and environment for which representative data is available. The model can then capture any dependencies present in the data. In current standard-defined aerial channels, for instance, the

distributions from which the angles of the multipath components are drawn do not depend on the distance; in contrast, and as intuition would have it, our model indicates, from the underlying data, a progressive narrowing of these distributions over distance.

In closing, we recall that, while the model has proved its ability to learn and to made interesting predictions driven by ray-tracing data, the ultimate objective is to drive it with empirical data. For this purpose, a measurement collection campaign is underway.

ACKNOWLEDGMENT

The handling by the Associate Editor, and the feedback provided by the anonymous reviewers, are gratefully acknowledged.

REFERENCES

- [1] W. Xia, S. Rangan, M. Mezzavilla, A. Lozano, G. Geraci, V. Semkin, and G. Loianno, "Millimeter wave channel modeling via generative neural networks," in *Proc. IEEE Globecom Workshops*, 2020, pp. 1–6.
- [2] G. Geraci, A. Garcia-Rodriguez, M. M. Azari, A. Lozano, M. Mezzavilla, S. Chatzinotas, Y. Chen, S. Rangan, and M. Di Renzo, "What will the future of UAV cellular communications be? A flight from 5G to 6G," *available as arXiv:2105.04842*, 2021.
- [3] Y. Zeng, I. Guvenc, R. Zhang, G. Geraci, and D. W. Matolak (Eds.), *UAV Communications for 5G and Beyond*. Wiley – IEEE Press, 2020.
- [4] W. Saad, M. Bennis, M. Mozaffari, and X. Lin, *Wireless Communications and Networking for Unmanned Aerial Vehicles*. Cambridge University Press, 2020.
- [5] 3GPP Technical Specification 22.125, "Technical specification group services and system aspects; unmanned aerial system (UAS) support in 3GPP; Stage 1; Release 17," Dec. 2019.
- [6] 3GPP Technical Specification 22.829, "Technical specification group services and system aspects; enhancements for unmanned aerial vehicles; Stage 1; Release 17," Sep. 2019.
- [7] G. Geraci, A. Garcia-Rodriguez, and X. Lin, "Preparing the ground for drone communications," in *IEEE ComSoc Technology News*, June 2019.
- [8] A. Fotouhi, H. Qiang, M. Ding, M. Hassan, L. Galati-Giordano, A. Garcia-Rodriguez, and J. Yuan, "Survey on UAV cellular communications: Practical aspects, standardization advancements, regulation, and security challenges," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 4, pp. 3417–3442, 2019.
- [9] M. Mozaffari, W. Saad, M. Bennis, Y. Nam, and M. Debbah, "A tutorial on UAVs for wireless networks: Applications, challenges, and open problems," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 3, pp. 2334–2360, third quarter 2019.
- [10] A. Garcia-Rodriguez, G. Geraci, D. López-Pérez, L. Galati Giordano, M. Ding, and E. Björnson, "The essential guide to realizing 5G-connected UAVs with massive MIMO," *IEEE Commun. Mag.*, vol. 57, no. 12, pp. 84–90, 2019.
- [11] G. Geraci, A. Garcia Rodriguez, L. Galati Giordano, D. López-Pérez, and E. Björnson, "Understanding UAV cellular communications: From existing networks to massive MIMO," *IEEE Access*, vol. 6, Nov. 2018.
- [12] X. Lin, R. Wiren, S. Euler, A. Sadam, H. Määtänen, S. Muruganathan, S. Gao, Y. . E. Wang, J. Kauppi, Z. Zou, and V. Yajnanarayana, "Mobile network-connected drones: Field trials, simulations, and design insights," *IEEE Veh. Tech. Mag.*, vol. 14, no. 3, pp. 115–125, 2019.

- [13] Y. Zeng, J. Lyu, and R. Zhang, "Cellular-connected UAV: Potentials, challenges and promising technologies," *IEEE Wireless Commun. Mag.*, vol. 26, no. 1, pp. 120–127, Feb. 2019.
- [14] M. M. Azari, F. Rosas, and S. Pollin, "Cellular connectivity for UAVs: Network modeling, performance analysis and design guidelines," *IEEE Trans. Wireless Commun.*, vol. 18, no. 7, pp. 3366–3381, July 2019.
- [15] H. C. Nguyen, R. Amorim, J. Wigard, I. Z. Kovács, T. B. Sørensen, and P. Mogensen, "How to ensure reliable connectivity for aerial vehicles over cellular networks," *IEEE Access*, vol. 6, pp. 12 304–12 317, Feb. 2018.
- [16] S. Singh, U. Bhattacharjee, E. Ozturk, I. Guvenc, H. Dai, M. Sichitiu, and A. Bhuyan, "Placement of mmWave base stations for serving urban drone corridors," in *Proc. IEEE VTC-Spring*, 2021, pp. 1–6.
- [17] S. Kang, M. Mezzavilla, A. Lozano, G. Geraci, W. Xia, S. Rangan, V. Semkin, and G. Loianno, "Millimeter-wave UAV coverage in urban environments," *arXiv preprint 2104.04600*, 2021.
- [18] 3GPP Technical Report 36.777, "Technical specification group radio access network; Study on enhanced LTE support for aerial vehicles (Release 15)," Dec. 2017.
- [19] M. Polese, L. Bertizzolo, L. Bonati, A. Gosain, and T. Melodia, "An Experimental mmWave Channel Model for UAV-to-UAV Communications," in *Accepted at the 4th ACM Workshop on Millimeter-wave Networks and Sensing Systems (mmNets'20)*, 2020.
- [20] H. Shakhathreh, W. Malkawi, A. Sawalmeh, M. Almutiry, and A. Alenezi, "Modeling ground-to-air path loss for millimeter wave UAV networks," *available as arXiv:2101.12024*, 2021.
- [21] S. Garcia Sanchez, S. Mohanti, D. Jaisinghani, and K. R. Chowdhury, "Millimeter-wave base stations in the sky: An experimental study of UAV-to-ground communications," *IEEE Transactions on Mobile Computing*, 2020.
- [22] R. Kovalchukov *et al.*, "Analyzing effects of directionality and random heights in drone-based mmWave communication," *IEEE Transactions on Vehicular Technology*, 2018.
- [23] M. T. Dabiri, H. Safi, S. Parsaeefard, and W. Saad, "Analytical channel models for millimeter wave UAV networks under hovering fluctuations," *available as arXiv:1905.01477*, 2019.
- [24] M. Gapeyenko, V. Petrov, D. Moltchanov, S. Andreev, N. Himayat, and Y. Koucheryavy, "Flexible and reliable UAV-assisted backhaul operation in 5G mmWave cellular networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 11, pp. 2486–2496, 2018.
- [25] 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, 2014.
- [26] T. S. Rappaport, R. W. Heath Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Pearson Education, 2014.
- [27] S. Wu, C.-X. Wang, M. Alwakeel, X. You *et al.*, "A general 3-D non-stationary 5G wireless channel model," *IEEE Trans. Commun.*, vol. 66, no. 7, pp. 3065–3078, 2017.
- [28] 3GPP Technical Report 38.901, "Study on channel model for frequencies from 0.5 to 100 GHz (Release 16)," Dec. 2019.
- [29] W. Khawaja, I. Guvenc, D. W. Matolak, U. Fiebig, and N. Schneckenburger, "A survey of air-to-ground propagation channel modeling for unmanned aerial vehicles," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 3, pp. 2361–2391, 2019.
- [30] R. Amorim, H. Nguyen, P. Mogensen, I. Z. Kovács, J. Wigard, and T. B. Sørensen, "Radio channel modeling for UAV communication over cellular networks," *IEEE Wireless Commun. Letters*, vol. 6, no. 4, pp. 514–517, 2017.
- [31] L. Cheng, Q. Zhu, C. Wang, W. Zhong, B. Hua, and S. Jiang, "Modeling and simulation for UAV air-to-ground mmWave channels," in *European Conf. Antennas and Propagation (EuCAP'20)*, 2020, pp. 1–5.
- [32] V. Semkin, S. Kang, J. Haarla, W. Xia, I. Huhtinen, G. Geraci, A. Lozano, G. Loianno, M. Mezzavilla, and S. Rangan, "Lightweight UAV-based measurement system for air-to-ground channels at 28 GHz," in *Proc. IEEE PIMRC*, 2021, pp. 1–6.

- [33] K. Stocker, B. Gschwendtner, and F. Landstorfer, "Neural network approach to prediction of terrestrial wave propagation for mobile radio," in *IEE Proceedings H (Microwaves, Antennas and Propagation)*, vol. 140, no. 4, 1993, pp. 315–320.
- [34] P.-R. Chang and W.-H. Yang, "Environment-adaptation mobile radio propagation prediction using radial basis function neural networks," *IEEE Trans. Veh. Techn.*, vol. 46, no. 1, pp. 155–160, 1997.
- [35] L. Bai, C.-X. Wang, J. Huang, Q. Xu, Y. Yang, G. Goussetis, J. Sun, and W. Zhang, "Predicting wireless mmWave massive MIMO channel characteristics using machine learning algorithms," *Wireless Commun. and Mobile Computing*, 2018.
- [36] J. Huang, C.-X. Wang, L. Bai, J. Sun, Y. Yang, J. Li, O. Tirkkonen, and M. Zhou, "A big data enabled channel model for 5G wireless communication systems," *IEEE Trans. Big Data*, vol. 6, no. 2, pp. 211–222, 2018.
- [37] X. Zhao, F. Du, S. Geng, Z. Fu, Z. Wang, Y. Zhang, Z. Zhou, L. Zhang, and L. Yang, "Playback of 5G and beyond measured MIMO channels by an ANN-based modeling and simulation framework," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 9, pp. 1945–1954, 2020.
- [38] E. Ostlin, H.-J. Zepernick, and H. Suzuki, "Macrocell path-loss prediction using artificial neural networks," *IEEE Trans. Veh. Techn.*, vol. 59, no. 6, pp. 2735–2747, 2010.
- [39] E. Dall'Anese, S.-J. Kim, and G. Giannakis, "Channel gain map tracking via distributed kriging," *IEEE Trans. Veh. Techn.*, vol. 60, no. 3, pp. 1205–1211, 2011.
- [40] L. Azpilicueta, M. Rawat, K. Rawat, F. Ghannouchi, and F. Falcone, "A ray launching-neural network approach for radio wave propagation analysis in complex indoor environments," *IEEE Trans. Antennas Propag.*, vol. 62, no. 5, pp. 2777–2786, 2014.
- [41] M. Kasparick, R. Cavalcante, S. Valentin, S. Stańczak, and M. Yukawa, "Kernel-based adaptive online reconstruction of coverage maps with side information," *IEEE Trans. Veh. Techn.*, vol. 65, no. 7, pp. 5461–5473, 2015.
- [42] G. Ferreira, L. Matos, and J. Silva, "Improvement of outdoor signal strength prediction in UHF band by artificial neural network," *IEEE Trans. Antennas Propag.*, vol. 64, no. 12, pp. 5404–5410, 2016.
- [43] D. Romero, S.-J. Kim, G. Giannakis, and R. López-Valcarce, "Learning power spectrum maps from quantized power measurements," *IEEE Trans. Signal Processing*, vol. 65, no. 10, pp. 2547–2560, 2017.
- [44] X. Ma, J. Zhang, Y. Zhang, and Z. Ma, "Data scheme-based wireless channel modeling method: motivation, principle and performance," *J. Commun. Inform. Netw.*, vol. 2, no. 3, pp. 41–51, 2017.
- [45] R. Nikbakht, A. Jonsson, and A. Lozano, "Dual-kernel online reconstruction of power maps," in *IEEE Global Commun. Conf. (GLOBECOM'18)*, 2018, pp. 1–5.
- [46] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [47] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [48] C. Doersch, "Tutorial on variational autoencoders," *arXiv preprint arXiv:1606.05908*, 2016.
- [49] Y. Yang, Y. Li, W. Zhang, F. Qin, P. Zhu, and C. Wang, "Generative-adversarial-network-based wireless channel modeling: Challenges and opportunities," *IEEE Commun. Magazine*, vol. 57, no. 3, pp. 22–27, 2019.
- [50] T. J. O'Shea, T. Roy, and N. West, "Approximating the void: Learning stochastic channel models from observation with variational generative adversarial networks," in *Int'l Conf. Computing, Netw. and Commun. (ICNC'19)*, 2019, pp. 681–686.
- [51] H. Ye, L. Liang, G. Y. Li, and B. Juang, "Deep learning-based end-to-end wireless communication systems with conditional GANs as unknown channels," *IEEE Trans. Wireless Commun.*, vol. 19, no. 5, pp. 3133–3143, 2020.
- [52] "Remcom," available on-line at <https://www.remcom.com/>.
- [53] A. Alkhateeb, "DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications," in *Proc. of Information Theory and Applications Workshop (ITA)*, San Diego, CA, Feb 2019, pp. 1–8.

- [54] V. Degli-Esposti, F. Fuschini, E. M. Vitucci, M. Barbiroli, M. Zoli, L. Tian, X. Yin, D. A. Dupleich, R. Müller, C. Schneider *et al.*, “Ray-tracing-based mm-wave beamforming assessment,” *IEEE Access*, vol. 2, pp. 1314–1325, 2014.
- [55] A. Al-Hourani, S. Kandeepan, and A. Jamalipour, “Modeling air-to-ground path loss for low altitude platforms in urban environments,” in *Proc. IEEE Globecom*, 2014, pp. 2898–2904.
- [56] Z. Yun and M. F. Iskander, “Ray tracing for radio propagation modeling: Principles and applications,” *IEEE Access*, vol. 3, pp. 1089–1100, 2015.
- [57] W. Khawaja, O. Ozdemir, and I. Guvenc, “UAV Air-to-Ground Channel Characterization for mmWave Systems,” in *Proc. IEEE VTC-Fall*, 2017.
- [58] F. Fuschini, S. Häfner, M. Zoli, R. Müller, E. Vitucci, D. Dupleich, M. Barbiroli, J. Luo, E. Schulz, V. Degli-Esposti *et al.*, “Analysis of in-room mm-wave propagation: Directional channel measurements and ray tracing simulations,” *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 38, no. 6, pp. 727–744, 2017.
- [59] “mmWave channel modeling git hub repository,” available on-line at <https://github.com/nyu-wireless/mmwchanmod>.
- [60] R. W. Heath Jr. and A. Lozano, *Foundations of MIMO Communication*. Cambridge University Press, 2018.
- [61] G. R. MacCartney, T. S. Rappaport, and S. Rangan, “Rapid fading due to human blockage in pedestrian crowds at 5g millimeter-wave frequencies,” in *IEEE GLOBECOM*. IEEE, 2017, pp. 1–7.
- [62] C. Slezak, V. Semkin, S. Andreev, Y. Koucheryavy, and S. Rangan, “Empirical effects of dynamic human-body blockage in 60 GHz communications,” *IEEE Commun. Magazine*, vol. 56, no. 12, pp. 60–66, 2018.
- [63] M. R. Akdeniz, Y. Liu, M. K. Samimi, S. Sun, S. Rangan, T. S. Rappaport, and E. Erkip, “Millimeter wave channel modeling and cellular capacity evaluation,” *IEEE journal on selected areas in communications*, vol. 32, no. 6, pp. 1164–1179, 2014.
- [64] A. A. Khuwaja, Y. Chen, N. Zhao, M. S. Alouini, and P. Dobbins, “A survey of channel modeling for UAV communications,” *IEEE Commun. Surveys & Tutorials*, vol. 20, no. 4, pp. 2804–2821, 2018.
- [65] R. Amorim, P. Mogensen, T. Sorensen, I. Z. Kovács, and J. Wigard, “Pathloss measurements and modeling for UAVs connected to cellular networks,” in *IEEE Veh. Techn. Conf. (VTC’17 Spring)*, 2017, pp. 1–6.
- [66] R. Amorim, H. Nguyen, J. Wigard, I. Z. Kovács, T. B. Sorensen, and P. Mogensen, “LTE radio measurements above urban rooftops for aerial communications,” in *IEEE Wireless Commun. Netw. Conf. (WCNC’18)*, 2018, pp. 1–6.
- [67] W. Xia, M. Polese, M. Mezzavilla, G. Loianno, S. Rangan, and M. Zorzi, “Millimeter Wave Remote UAV Control and Communications for Public Safety Scenarios,” in *16th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, 2019.
- [68] A. Lozano and N. Jindal, “Are yesterday’s information-theoretic fading models and performance metrics adequate for the analysis of today’s wireless systems?” *IEEE Commun. Magazine*, vol. 50, no. 11, pp. 210–217, 2012.
- [69] Z. Chen, H. Gao, D. Leenaerts, D. Milosevic, and P. Baltus, “A 29–37 GHz BiCMOS Low-Noise Amplifier with 28.5 dB Peak Gain and 3.1–4.1 dB NF,” in *Proc. IEEE RFIC*, 2018, pp. 288–291.
- [70] R. Garg and A. S. Natarajan, “A 28-GHz Low-Power Phased-Array Receiver Front-End With 360° RTPS Phase Shift Range,” *IEEE Trans. Microw. Theory Tech.*, vol. 65, no. 11, pp. 4703–4714, 2017.
- [71] L. Rüschendorf, “The Wasserstein distance and approximation theorems,” *Probability Theory and Related Fields*, vol. 70, no. 1, pp. 117–129, 1985.
- [72] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, “Improved training of Wasserstein GANs,” *arXiv preprint arXiv:1704.00028*, 2017.
- [73] S. Vallender, “Calculation of the wasserstein distance between probability distributions on the line,” *Theory of Probability & Its Applications*, vol. 18, no. 4, pp. 784–786, 1974.