Outdoor Millimeter-Wave Picocell Placement using Drone-based Surveying and Machine Learning

Ian Chandler McDowell, Rahul Bulusu, Hem Regmi, Sanjib Sur
Computer Science and Engineering; University of South Carolina, Columbia, USA
{mcdoweli, rbulusu, hregmi}@email.sc.edu, sur@cse.sc.edu

***Abstract*—Millimeter-Wave (mmWave) networks rely on carefully placed small base stations called "picocells" for optimal network performance. However, the process of conducting site surveys to identify suitable picocell locations is both expensive and time-consuming. The current low-cost approaches for indoor surveying are often unsuitable for outdoor environments due to the presence of various environmental factors. To address this issue, we present *Theia*, a drone-based system that predicts outdoor mmWave Signal Reflection Profiles (SRPs) and facilitates picocell placement for optimal network coverage. The drone platform integrates optical systems and a mmWave transceiver to collect depth images and mmWave SRPs of the environment. These datasets are fed into a machine learning model that maps the depth data to SRPs, allowing SRPs to be predicted at previously unseen parts of the environment. *Theia* then leverages these predictions to identify optimal picocell locations that maximize network coverage and minimize link outages. We evaluate *Theia* in three large-scale outdoor environments and demonstrate that the proposed design can generalize the deployment method with a little refinement of the model.**

***Index Terms*—Drone; Millimeter-Wave; Outdoor Picocell Deployment; Deep Learning.**

I. Introduction

5G technology has been rapidly proliferating over the last few years, and it is expected to support the connectivity of massive numbers of smart devices worldwide [1]. As the use of 5G increases, the need for proper network infrastructure deployment becomes ever more important [2]. 5G uses millimeter-wave (mmWave) as its major wireless technological component and provides connectivity with picocells. Picocells are small, short-range base stations with the ability to produce higher data rates than traditional networks. However, there are multiple challenges in finding optimal locations for deploying these picocells: Picocells are easily obstructed by the environment in Line-of-Sight (LoS) paths, and hence they need to rely heavily on Non-Line-of-Sight (NLoS) paths. This is especially apparent in urban settings where NLoS paths are much more frequent due to dense buildings and infrastructure [3]. However, due to the high operating frequencies of picocells, many objects in the outdoor environment do not reflect the signals strongly, and therefore availability of NLoS paths becomes a challenge. So, deployers must find the locations where strong NLoS paths are available most of the time during LoS path unavailability. Existing approaches to finding such locations rely mostly on thorough site surveys or propagation models. They both aim to measure or model the mmWave Signal Reflection Profile (SRP) of different locations, and based on the estimated SRP, they find the picocell locations that will likely achieve high signal strength through NLoS paths.

However, site surveys are expensive [4], and propagation modeling and simulations are often unable to faithfully represent the environment at higher wireless frequencies. Besides, manual site surveying of a large-scale outdoor environment is not only time-consuming but also oftentimes infeasible. Past works have aimed to reduce survey time to find optimal picocell locations [5], [6], but they are limited to small-scale indoor environments, and hence do not address the difficulties present in outdoor environments such as area-span of the environment, foliage, number and height of buildings, pedestrian and automotive traffic, *etc.*

In this work, we propose Theia, which aims for easier surveying of outdoor environments and enables the deployers to identify near-optimal picocell locations for deployment.[1] *Theia* aims to identify the SRPs of different locations in a way that allows network deployers to survey an environment faster while still maintaining the accuracy of traditional surveys. Since SRPs include the strength of NLoS paths, the network can be more effective in dynamic outdoor environments by finding locations that maximize the SRP. Instead of measuring the SRP from all locations, *Theia* leverages the visual information of the environment and models the complex relationship between the visual to mmWave SRP. Since it is intuitive that similar-looking objects likely reflect mmWave signals similarly (e.g. varying types of foliage, cars, buildings, etc.), the model could predict the SRP from unseen viewpoints in the environment, requiring less tedious measurement from the deployer. To design the model, *Theia* leverages the advancements of deep learning frameworks and builds a customized deep convolutional neural network that can predict the mmWave SRPs from visual data of the environment. These predicted SRPs are then used to find the picocell locations that will maximize the availability of the NLoS paths.

We implement and evaluate *Theia* with a custom-built platform for data collection. The platform uses a DJI Matrice 100 Drone [7] with a DJI Guidance System [8] to collect the depth images and poses of the system, and a mounted 77 GHz mmWave transceiver [9] is used for collecting SRPs. Our experimental results across 6 outdoor environments over a period of 4 months with 44 GB of collected data (∼144,000

[1] *Theia* is the Greek goddess of sight and vision which is analagous to our goal of surveying large outdoor environments.

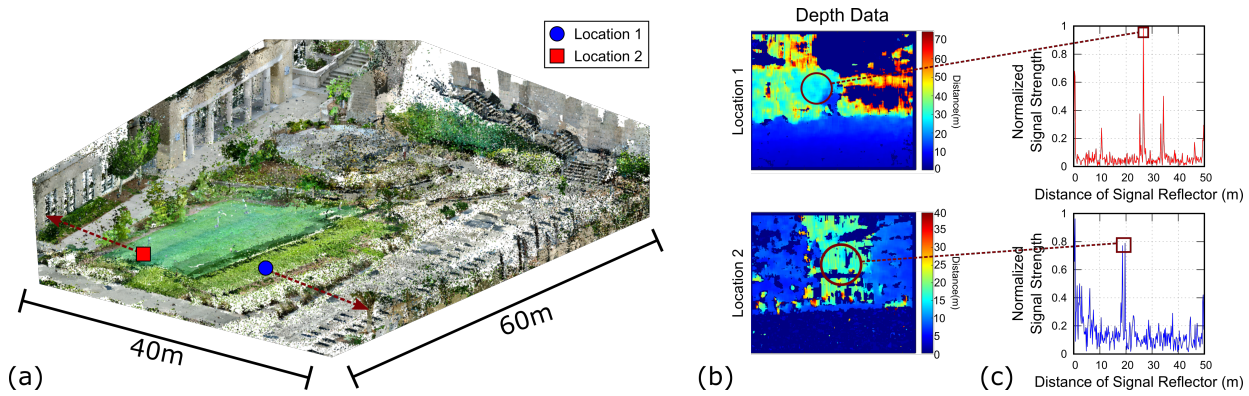

Figure 1: (a) Point Cloud Data (PCD) of an outdoor environment. (b–c) Depth images and mmWave SRPs from 2 locations in the environment.

samples), show that *Theia* achieves a median error of 3.73 dB in SRP prediction with a model trained for a unique environment. Furthermore, *Theia* only requires 5 minutes of survey data in a new environment with different amounts and types of objects to reach similar prediction performance. For picocell deployment, *Theia* can achieve a 2.76× reduction in link outage likelihood in multiple environments compared to random and common-sense deployment strategies.

In summary, we have the following contributions: (1) We build a drone system for automatically collecting and synchronizing visual data and mmWave SRPs from outdoors. (2) We design a custom deep learning model for predicting SRPs in the environment, which allows the network deployer to find near-optimal picocell locations to improve the performance and reliability of an outdoor mmWave network. Our results demonstrate that *Theia* can generalize in different outdoor environments with little fine-tuning.

II. Background and Challenges

A. Picocell Fundamentals and SRP

MmWave networks rely on picocells to provide wireless connectivity within a limited range of approximately 10-12 meters [10]. Due to the presence of obstacles on the LoS path, NLoS paths are frequently used, underscoring the importance of determining the SRP resulting from the reflectors present in the environment. The SRP can be obtained by sending a signal at varying frequencies in the given environment. By determining the time taken for the reflected signal to return, different distances can be estimated. To accomplish this, *Theia* utilizes a mmWave transceiver to send an extensive full-bandwidth signal into the environment, containing multiple reflecting points. The total reflected signal received when the signal bounces back is then measured by summing the time delays of the signal. Applying Fast Fourier Transform (FFT) to this signal enables us to obtain various signal strengths from different object distances [11]. *Theia*'s transceiver is equipped with multiple transmit and receive antennas, allowing us to obtain a multi-channel SRP by emitting a signal from each transmit antenna and measuring the reflected signal at each receive antenna. Since each antenna is positioned differently, each signal differs and has a virtual pose of origin based on the position of the transmit and receive antennas. Figure 1 shows an example outdoor environment and the visual data and mmWave SRPs from two different locations.

B. Propagation Survey Challenges

To ensure optimal placement of picocells in outdoor environments, precise planning is crucial. Failure to locate picocells in the right position can result in various issues, including capacity loss and obstruction by environmental objects. Existing approaches to finding such locations traditionally relied on two main approaches: site surveys and propagation models. However, both of these methods have significant drawbacks. Site surveys are expensive and time-consuming [4], making them impractical for outdoors. Propagation modeling and simulations are often inaccurate, especially at higher wireless frequencies, where the environmental reflection is more complex. Although recent surveying methods reduce time and cost, they are limited to indoor environments [5], [6]. These methods are inadequate for outdoor environments because they do not consider environmental factors such as building height, population density, and limited area coverage, which can hinder their effectiveness.

III. *Theia* Design

A. Overview

Theia allows for easy SRP prediction in outdoor environments, which enables network deployers to achieve optimal picocell placement with predictable performance. To get the SRP predictions, *first*, the deployer uses a DJI Matrice 100 Drone with a DJI Guidance System [7], [8], which is equipped with a depth sensor, to collect pose data and depth images of an outdoor environment by flying the drone. *Second*, as the deployer is flying the drone, a co-located mmWave transceiver measures the SRPs of the environment by steering the mmWave beam. *Finally*, *Theia* uses depth data and SRPs to learn the relationship between objects in the environment and their reflections, which facilitates the prediction of full SRPs for the environment.

To achieve this, *Theia* uses a Convolutional Neural Network to map the pose and depth image to the reflection using supervised learning. Figure 2 shows an overview of the

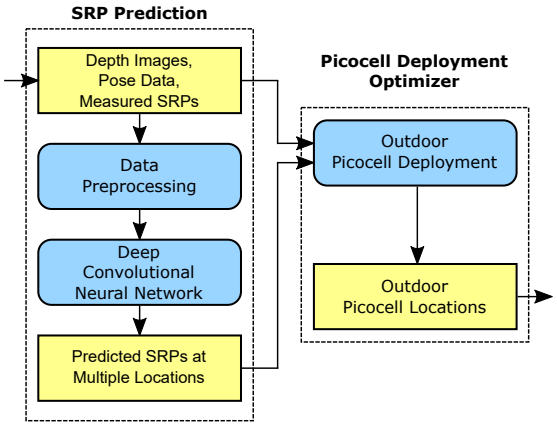

Figure 2: Overview of the *Theia* system.

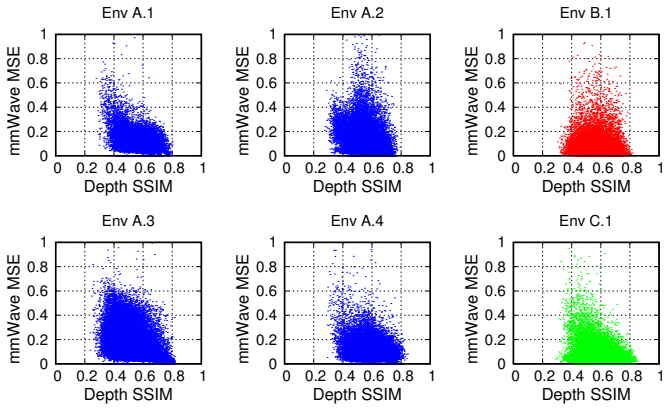

Figure 3: Visual data and SRP relationship across 6 environments.

Theia system. The depth images and drone poses are first synchronized with the measured SRPs collected at different positions during flight. The transceiver's beam pattern is then masked over the depth image and paired with the measured pose and SRP, to train the prediction network. The network uses thousands of collected data pairs to learn the relationship between the visual data and measured SRP and harnesses this to predict the SRPs at different environmental positions. This model is generalizable to other outdoor environments with similar objects. Finally, the predicted SRPs can be used for finding picocell placements with predictable performance.

B. Visual Data to SRP Relationship

Before developing a model for SRP prediction from visual data, we first investigate the relationship between visual data and SRPs to test the hypothesis: *similar objects produce similar mmWave reflection profiles* (e.g. varying types of foliage, cars, buildings, etc.), so that the model can learn SRP from one part of the environment and predict SRP in another similar-looking part. To test this hypothesis, we compute the Structural Similarity Index Measure (SSIM) [12] between two depth images to represent visual similarity and the Mean Squared Error (MSE) of the corresponding reflections to represent SRP similarity. We conduct this analysis on data pairs from six environments across three outdoor spaces, generating ∼144K data pairs. For details on each environment see Table I. Figure 3 shows a trend between visual data and SRPs, but the relationship is complex and nonlinear. For example, in some environments, like A.2 and B.1, the relationships look similar, but in other environments, they are quite different. A regression model would not capture this trend, and different patterns across environments make a single model non-generalizable. To capture such complex relationship, we adopt a deep learning approach in a generalizable manner, allowing us to predict SRPs from angles and points we did not observe directly, saving time during surveying.

C. SRP Prediction using DCNN

To learn the mapping between depth images and reflections, *Theia* utilizes a Deep Convolutional Neural Network (DCNN). The problem of predicting signal strength in an environment

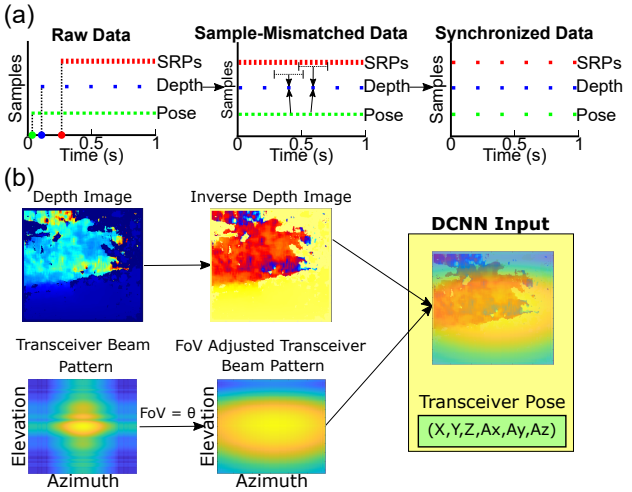

Figure 4: (a) Data synchronization for SRP, pose, and depth image. (b) Constructing DCNN input (θ is 60°).

is in the same vein as the image super-resolution problem [13], [14], [15], [16], [17], with the measured reflections corresponding to low-resolution images where the reflections are measured from a limited number of positions and sent through a DCNN, with the positions acting as labels in order to reconstruct the higher resolution images. To this end, we feed our depth image as input to the DCNN with the SRP collected at the same pose being the ground truth. Furthermore, we include the antenna pose in the network since the SRP can be affected significantly by how the deployer is holding the device. Finally, we use the MSE between the predicted and ground truth SRPs as the loss function for the network.

1) Data Preprocessing: To facilitate proper training, *Theia* first synchronizes the raw data, then removes noise and unnecessary information (see Figure 4). Since a tight hardware-level synchronization does not yet exist between the drone, transceiver, and depth imager of our data collection platform, we implement a software-level synchronization. Since the systems do not start simultaneously, and the flight pose and depth imager have a lower sampling rate than the transceiver, we must process the data to align it. In particular, the flight pose has a rate of 20 samples per second (sps), the depth

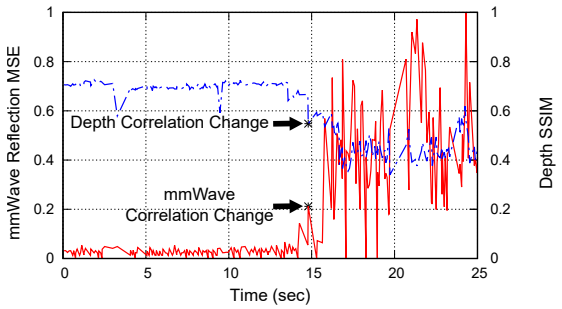

Figure 5: Back-to-back correlation of mmWave and depth data samples are used to confirm proper synchronization between the mmWave reflections and depth images.

imager has a rate of about 6 sps, and the transceiver has a rate of 25 sps. To rectify the differences in start time, we have each system record the start timestamps, and we use the difference in start times to align the data. Once the flight, depth, and transceiver data are aligned, we use interpolation and decimation to match the rates across all three systems. We have chosen to specifically use piecewise-cubic interpolation and median decimation because they produce good results.

Since it is infeasible to determine a time cutoff when transceiver samples are related to each depth image due to the significant difference in sampling rates, we use an overlapping time window approach to get the transceiver samples closest in time to the depth image and take the median of those samples as the corresponding sample for the depth image, as shown in Figure 4(a). Also, back-to-back samples will be similar during the stable 10 seconds at the beginning of data collection. To verify whether the data is synchronized, we use the intuition that back-to-back samples will be similar until the drone starts moving. We first correlate the back-to-back samples using Mean Squared Error (MSE) for the transceiver and drone pose data and Structural Similarity Index Measure (SSIM) [12] for the depth data. Then, by looking at where the correlation changes significantly, we can identify the point in the time axis where the drone started to move, as seen in Figure 5. We do not use the back-to-back correlation to synchronize as the threshold that must be reached to determine if the drone has started varies from environment to environment and thus would require tuning for each new data collection. Figure 6 shows two examples of the depth and 12 corresponding mmWave SRPs from each virtual channel. Each example is from a different pose in two different environments after synchronizing the data.

Now that the data is properly synchronized, we prepare the data for input into the DCNN model. Because of the transceiver's specific signal emission (beam) pattern [18], [19], not all of the information in the depth image will contribute equally to the measured SRPs. To account for this, we first get the Inverse Depth Image (IDI) to explicitly tell the DCNN to prioritize closer objects when learning what contributes to strong reflections. The IDI is the inverse depth value at each pixel. We then need to include the mmWave transceiver's beam pattern [20], a 2D matrix with the normalized transmit and receive power for different elevation and azimuth angles. However, since the depth sensor's Field of View (FoV) is much smaller, the beam pattern is limited to the FoV that overlaps with the depth sensor's. Then, via dot multiplication, we mask the IDI with the FoV-adjusted beam pattern, getting a *Masked Inverse Depth Image* (MIDI). The MIDI encodes information both about the reflectivity of objects and takes into account the properties of the transceiver. The MIDI is then associated with a pose and SRP, which forms the input and output data for the DCNN, as shown in Figure 4(b).

2) DCNN Model: Although preprocessing allows us to remove unnecessary features and prioritize the important ones, learning a mapping from the MIDI and pose to the SRP is still challenging as the relationship is complex and non-linear (see Figure 3). Machine learning models have been used to tackle such non-linear tasks, and with recent developments in wireless use-case deep learning [5], [21], [22], [23], we find this is a suitable method to build a model for *Theia*. Since the SRP is just a vector of signal strength values, we build a DCNN model that uses a MIDI and pose as inputs and produces the associated SRP as the output. Now with the use of Convolutional layers [24], *Theia* can extract features from the MIDI, with its image-like structures, and feed them into Fully Connected (FC) layers for SRP prediction [25]. With so many Convolution and FC layers, the DCNN will have millions of parameters and introduce significant memory and computation constraints for *Theia* [26]. For this purpose, we select a model with lower memory and computational requirements that does not sacrifice accuracy.

To find a model with lower memory use and shorter training time, as well as a better SRP prediction accuracy, we evaluated multiple popular models as Convolution layers: *VGG16*[27], *EfficientNet*[28], *MobileNetV2*[29], *MobileNetV3-Large*, and *MobileNetV3-Small*[30] and ranked them by their MSE loss on test samples. We choose *MobileNetV3- Large* as it has the best performance of all the models tested and has memory usage and training time only second to *MobileNetV3-Small*.

MobileNetV3 is designed for mobile devices with lower memory and computational resources and builds off of its predecessor *MobileNetV2*, which uses depthwise separable convolutions, allowing lightweight filters to replace full convolution layers, significantly reducing the computational cost. *MobileNetV2* also replaces standard residual blocks with inverted blocks to significantly reduce the required tensors for storing intermediate results, which uses memory more efficiently. *MobileNetV3* then builds off its predecessor through the use of neural architecture search [31], [32], which utilizes a recurrent neural network to generate more computationally efficient network architectures.

Since the standard input of *MobileNetV3-Large* is a 3-channel 224×224 RGB image, we customize it for *Theia*. The MIDI in *Theia* is a monochrome 230×246 image, so we concatenate the MIDI across the 3 channels to generate a 230×246×3 input that we supply to the *MobileNetV3-Large*.

Since *MobileNetV3* is designed and trained on the ImageNet

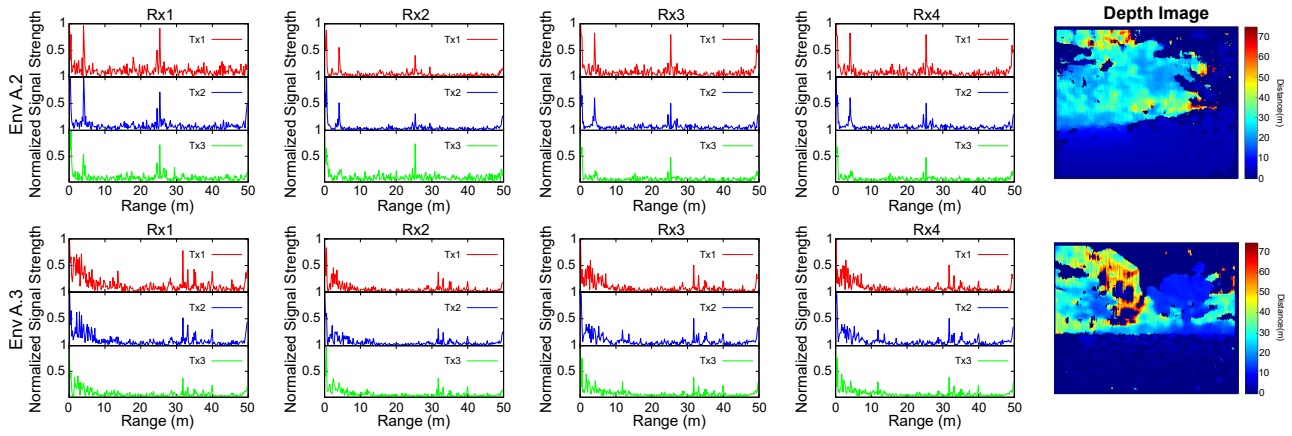

Figure 6: Synchronized data from two environments: SRPs across 12 transmit-receive channels; Depth images.

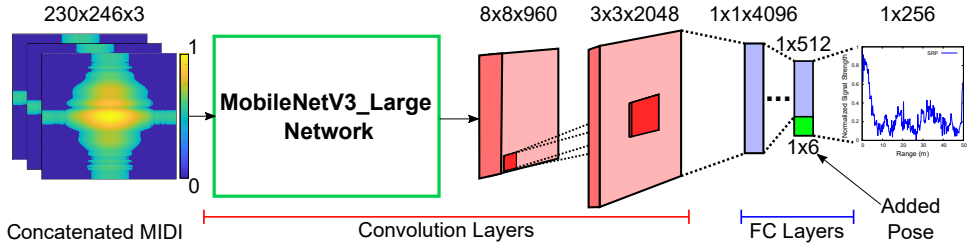

Figure 7: Customized DCNN model for *Theia*.

dataset [33] for image classification, it will not achieve *Theia*'s goal of predicting SRPs as that is a regression task. For this reason, we trim *MobileNetV3* to its convolution layers for feature extraction, then feed these to the customized FC layers for regression. We also provide the pose in the second to last FC layer since mmWave SRP also depends on the transceiver's location and rotation (*i.e.*, pose). Thus, including the pose directly in the DCNN FC layer increases the network's generalizability in robust environments. Figure 7 shows the DCNN model architecture in *Theia*.

D. Picocells Deployment from the Estimated SRPs

Due to the narrow beamwidth of the mmWave, even a minor blockage, such as hand movement, building structure, or another person passing by can break the link between the picocell and the client. Frequent link drops lead to a drop in the effective data throughput, increase the network latency and a need for multiple beam alignments. Even though it is always preferred to have a LoS path because it provides higher SNR, we can increase network reliability by considering NLoS paths, where a picocell can steer its beam toward an available strong reflector and establish the link to provide seamless coverage. SRPs provide information about strong reflectors in the environment, and we can use them to place picocells to maximize the use of the strong reflectors.

Then, to identify the picocell locations, we adopt a placement algorithm similar to that proposed in [5], which employs a Ray-tracing approach [34] to simulate separate transmitter and receiver and incorporate realistic reflections using the predicted SRPs. We derive three placement strategies, based on different coverage objectives and apply them to determine the picocell locations for *Theia*. The first strategy, "average," aims to provide higher mean throughput across all clients in the environment. The second strategy, "variance," aims to achieve nearly equal signal strengths across different parts of the environment to improve fairness. The third strategy, "link-outage," involves placing picocell locations to ensure a certain level of signal strength across most of the environment to minimize the likelihood of outages.

IV. Experimental Setup and Implementation

Hardware Setup: Current commodity 5G mmWave devices such as phones do not provide raw SRPs, so we opt to create a custom data collection platform to get the SRPs and to evaluate *Theia*. Our data collection platform comprises three systems that work in parallel to survey outdoor areas. (a) The flight system uses DJI's Onboard SDK [35] along with a USB to UART connection between the drone and laptop to control a DJI Matrice 100 Drone [7] along a programmable flight path and saves the pose data. (b) The guidance system uses DJI's Guidance SDK [36] to collect readings from the DJI Guidance System [8], including stereo grayscale visual and depth (RGB-D) images as well as flight odometry samples and saves them over USB to the laptop. (c) We mount a TI IWR1443BOOST 76-81GHz mmWave transceiver and a DCA1000EVM capture card [9], [37] on top of the drone that is controlled and powered over USB and saves data over an ethernet cable. The transceiver's 3 transmit and 4 receive antennas can collect SRPs across 12 virtual channels continuously. We placed the transceiver so that the drone's propellors would not damage

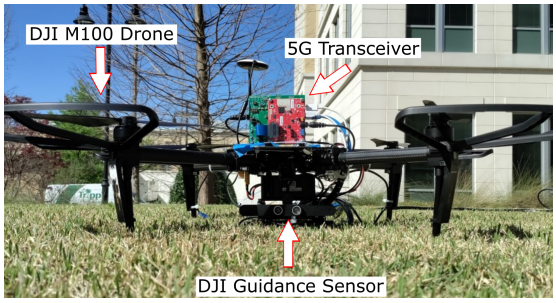

Figure 8: Data collection platform with a 76–81 GHz mmWave transceiver mounted on a DJI Matrice 100 drone equipped with DJI Guidance depth and grayscale visual sensors.

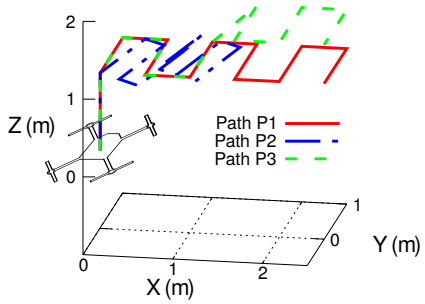

Figure 9: Different drone waypoint paths used for data collection: (P1) Zigzag with static altitude; (P2) Angled zigzag with altitude change halfway through a pass; and (P3) Zigzag with altitude change halfway through a pass. For the first half of the pass, the drone performs a 30° clockwise (CW) yaw rotation at each corner. For the other half of the pass, each rotation is counter-clockwise (CCW).

it; however, this resulted in a significant reflection from the drone around 0 m in the measured SRPs.

Since we are in the process of acquiring FAA certification, and due to difficulties in sending mass amounts of data wirelessly, we have opted to keep a wired, 2 meter connection between the drone and laptop. The data collected is within 5 meters of the ground and ensures the data is representative of the ground-based users that *Theia* intends to optimize for. The cables are bundled together and managed in a way that prevents damaging them with the drone's propellers and reduces the likelihood of them being unplugged mid-flight which would cause data loss.

The transceiver uses a 764 MHz bandwidth with a carrier frequency of 77.38 GHz. We use the following parameters for our SRP measurements: start frequency, 77GHz; frequency ramp slope, 29.98 MHz/μS; number of ADC samples, 256; ADC sampling rate, 10 Msps; sweep duration, 60 μs; pulse repetition rate, 25 Hz; and maximum receive antenna gain, 30 dBi. The measured SRP for each virtual channel is a 256 element vector. Each SRP element has a resolution of $\sim$0.1961 m [38]; thus, the device can collect reflections up to $\sim$49.73 m. Since the difference between the transceiver and drone positions is static during every data collection, this offset can be used for calibrating the transceiver's pose.

Real Data Collection: Using a laptop that is connected to the drone, a Matlab program runs to start the flight and guidance systems to ensure that there is no issue with the drone and its sensors. Next, the transceiver goes through its startup processes. When all the systems are ready to go, the drone waits 10 seconds, takes off, confirms sane output from the pose sensor, and follows the programmed flight pattern while the guidance system and transceiver record their measurements. Despite using commercial-off-the-shelf components, a hardware-level synchronization does not yet exist, so these systems do not start at the same time and they have different sampling rates. However, recall that we address these synchronization issues in preprocessing (Section III-C). We have chosen to use a combination of zigzag patterns, which consists of a forward and backward pass. The forward pass is shown in Figure 9, and the backward pass is the forward pass in reverse. With a selected path and varying yaw angles, we survey the area for a 10-minute duration, gathering on average 4800 unique transceiver poses.

During data collection, one person is holding the laptop and cable leash to the drone to ensure everything is running properly and the drone does not fly off, while a second person has the manual controller to keep an eye on battery life, GPS signal, and manually override the drone, should it veer dangerously off course. To maximize the data collected and overcome the limited battery life of the drone platform, we follow a battery cycling procedure where multiple drone batteries are kept charged, and after each data collection, a new one is placed into the drone. This deals with the problem of limited drone battery life; however, the laptop is still limited in its own battery life, so when it has reached its limit, we take the system to a place to charge and, while doing so, take the time to backup the new data and clean out the old data to allow for more frequent data collection.

Even with set waypoints, weather and GPS conditions significantly affect the drone's ability to follow them, so it is imperative to check that conditions are clear enough for the environment in which data is being collected. For example, in a larger, less trafficked area, higher winds and weaker GPS may be acceptable, but smaller environments with a lot of close-by obstacles would need slower wind speeds and higher GPS activity to avoid crashing. Using tools such as [39] saves time and system damage is avoided.

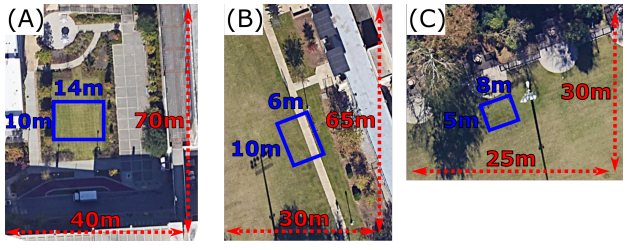

Figure 10: Top down view of each outdoor space with rectangles indicating the area where the drone flew and arrows indicating the area the drone observed.

We collect datasets across 6 environments in 3 different outdoor spaces over a period of 4 months. Our spaces are primarily used as a courtyard and sports field areas with different sizes and types of objects. For a detailed breakdown

Table I: Different properties of data collection in the 6 outdoor environments.

Environment	Drone Path	Base Yaw Angle	Purpose	Elements of Environment
A.1	P1	0°	Courtyard and Patio	Trees, patio tables, patio seating, stairs, handrails
A.2	P1	90° CW	Courtyard	Trees, benches, glass windows
A.3	P2	45° CCW	Courtyard and Patio	Trees, patio tables, patio seating, stairs, glass windows and doors
A.4	P3	0°	Courtyard and Patio	Trees, patio tables, patio seating, glass windows and doors
B.1	P3	0°	Sports Field	Trees, bushes, benches, fences, glass windows and doors
C.1	P1	180° CW	Sports Field Lounging Area	Trees, bushes, picnic table, fences

of each environment, see Table I, and a top-down view of each outdoor space can be seen in Figure 10. Each environment in the table takes place in the outdoor space of the corresponding letter (*i.e.*, each A.# takes place in outdoor space A but with different paths and starting angle. Since the data is collected during regular business hours, there could be background noise from pedestrians. In total, we have collected and processed over 44 GB of data with ∼144,000 data samples.

Neural Network Training: *Theia*'s SRP prediction model is trained and tested using MSE loss. While training, the network may continue for up to 1000 epochs; however, if learning stagnates, causing no improvement for 20 epochs in a row, the network will halt training. We tested multiple different optimizers and learning rates for our DCNN models by following [40] and saw that "NAdam" performed the best with a learning rate of 0.005. We also tested the MSE, L1, and SmoothL1 loss functions and found that they all produce similar results, so we chose MSE as the loss function for our model. The final models are designed and implemented using the Python programming language and PyTorch 1.10.2 [41] package on a server with AMD EPYC CPU @ 2.8GHz, 264GB RAM with Nvidia's RTX A6000 GPU [42] to decrease training time to about an hour for each model.

V. Experimental Results

Evaluation Summary: *Theia*'s DCNN model predicts SRP with a median error of 3.73 dB and needs only 5 minutes of surveying data to fine-tune to new environments. *Theia* performs well regardless of the number of mmWave channels used in training and the path variability of the drone while surveying the environment. When trained in completely new outdoor environments, *Theia* performs identically. *Theia*'s picocell deployment performs close to optimal with its "link-outage" strategy reducing the area without link by ∼2.76× compared to the Random and Common-Sense methods.

A. SRP Prediction

DCNN Model: To evaluate the performance of *Theia*'s DCNN model, we use datasets from outdoor environment A, only using the samples corresponding to one of the virtual channels. We preprocess the data to generate input pairs <MIDI, Pose, SRP>, then randomly choose samples with a 90:10 ratio of training to testing. We then shuffle the training samples and feed them into the model. We train the model using MSE loss, and once training is complete, we input MIDI and pose pairs from the testing samples to the model to predict the SRPs. We then calculate the absolute error between the predicted and ground truth SRPs for the 2145 test samples. Figure 12(a) shows the SRP error distribution with a median prediction error of 3.73 dB and 90th percentile error of 9.32 dB.

Multi-Channel Surveying: Since a deployer's device will have a fixed number of channels, we evaluate the performance of *Theia* when trained with a varying number of virtual channels. Specifically, the number of channels we chose to evaluate the performance for *Theia* are as follows: 1, 2, 3, 4, 6, 9, and 12. For each of the channels, we evaluated them with the same metrics. As shown in Figure 11(a) all the channels had similar results. Notably, 6 channels had the lowest median error of 3.571 dB and had a 90^{th} percentile error of 8.65 dB, and 4 channels had the highest median error of 3.848 dB and the highest 90^{th} percentile error of 9.54 dB. As the performances of each number of virtual channels tested are not significantly varying, *deployers won't have to worry about how many virtual channels a mmWave device has, as it has minimal effect on the performance*.

Drone Path Variability: To explore the importance of data collection pose variability, we evaluate *Theia* using data collected with low, medium, and high pose variability (drone paths P1, P2, and P3 in Figure 9, respectively). For this, we used data collected from environments A.1, A.3, and A.4, respectively. We evaluated each pose variance with the same metrics. As shown in Figure 11(b), we can see that the pose variability gave a similar SRP prediction error. From these results, since the errors were within 0.2 dB on both metrics, the path used to survey an environment appears to have little effect on SRP prediction. While further testing will need to be done to understand the impact of pose variability better, *these results show that the deployer will likely not need to worry about pose variability during surveying*.

Survey Time Requirement for Fine-Tuning: To determine the effect of fine-tuning in a new environment, we evaluate *Theia* by testing its performance with varying amounts of survey time. We train the model on data from environment A.1, then train it further with data from environment A.2. The number of training samples from A.2 is based on the amount of survey time being used for fine-tuning. We then test the model on data from environment A.2. We first test *Theia* using 0 minutes of survey time to see the result of no fine-tuning. We then test with 1, 5, 10, 15, 20, and 25 minutes of fine-tuning. We evaluate each survey time with the same metrics. Our base test, 0 minutes of survey time, had a higher median error of 6.341 dB and a 90th percentile error of 13.35 dB, which is intuited with no fine-tuning. With 1 min of fine tuning, we see

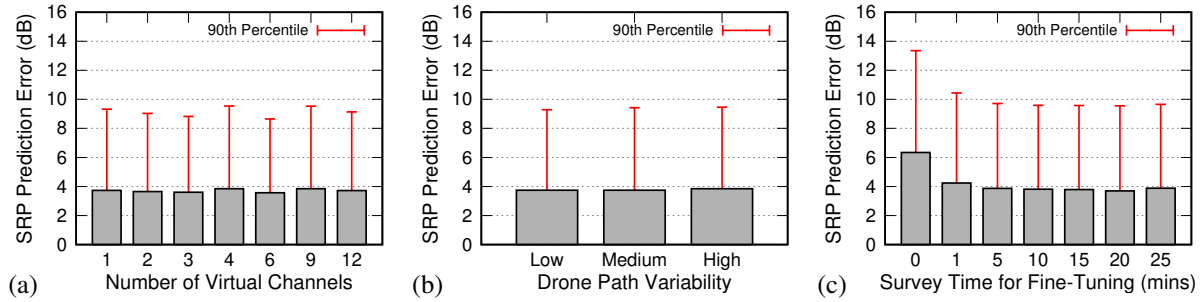

Figure 11: (a) Effect of the number of Rx-Tx combinations on model performance. (b) Comparing model performance on data with low, medium, and high pose variability. (c) Effect of survey time on model fine-tuning.

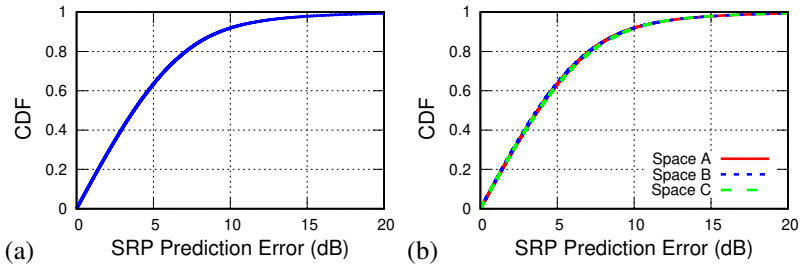

Figure 12: (a) CDF of SRP prediction error in Env A. (b) Performance of *Theia* when trained in different environments and tested on unseen parts of the environment.

that we already reduce the median error to 4.239 dB and the 90th percentile error to 10.44 dB, and with fining-tuning for 5 minutes, we get equivalent performance to when the model is tested on the same environment as training. 5, 10, 15, 20, and 25 minutes of fine-tuning all had similar performance results, as shown in Figure 11(c), meaning that there is limited gain in continuing to fine tune the model with longer survey times. These results indicate that *Theia* can adjust to environments that have never been surveyed. *Moreover, the deployer can save time collecting new samples.*

Generalizability to New Environments: We now evaluate *Theia*'s performance in different outdoor environments. We train models for each outdoor space with samples from environments A.1, B.1, and C.1, respectively. We evaluated the performance with the same metrics. Despite the differences between each of the environments, we see in Figure 12(b) that *Theia* is still able to predict the SRPs with all errors within $\sim$0.1 dB of each other. *So, Theia is able to perform well when trained in new environments without changing the model's parameters.*

Picocell Deployment Locations from Predicted SRP: With accurate SRP values, we can build a Ray-tracing method that will simulate a separate transmitter and receiver; and consider realistic mmWave reflections. We then use the mean and standard deviation from *Theia*'s prediction error as noise onto the simulated SRPs. This allows us to use realistic errors in our placements and enable comparison of *Theia* with the Optimal method in which the SRPs of the environment are known. We then follow the 3 deployment strategies, "average," "variation," and "link-outage" [5], to evaluate their performances. We also simulate three additional deployment methods: "Random," which chooses random positions in the environment for placement; "Common-Sense," which places the picocells at the corner locations; and "Optimal," which assumes there is no prediction error on SRPs. Figure 13(a–c) shows the placement of 64 picocells predicted using the "link-outage" strategy for each environment, and compares the results to the "Optimal," "Random," and "Common-sense" placements. The number of picocells is based on the environment size and estimated number of users, however a robust way to determine this number is something we plan to explore in the future. The results show a significant overlap between *Theia* and Optimal, meaning that *Theia* will likely perform comparably to an optimal deployment. Figure 14(a–c) shows the results for each strategy in environment A as each environment performed similarly for the respective strategies. Figure 14(a) shows all placement strategies as within 5 dB on the median. Figure 14(b) shows that *Theia* will limit the SRP variation to 0.8 dB whereas the Random and Common-Sense deployments can get up to 1.75 dB. Figure 14(c) shows that *Theia* reduces the area without links by $\sim$2.76$\times$ compared to Random and Common-Sense. *These results show the importance of accurate SRP prediction for deploying picocells in outdoor areas.*

VI. Related Works

MmWave Signal Reflection Prediction: Accurately obtaining signal strength maps is crucial for planning and operating networks, but it is often expensive and prone to errors for both carrier and crowdsourcing companies [43]. While the Friis path loss model [44] is effective for lower frequencies, it fails at high frequencies due to the small wavelengths that prevent diffraction when hitting obstacles [45]. Propagation simulators [46], [47], [48], [49] are commonly used for low-

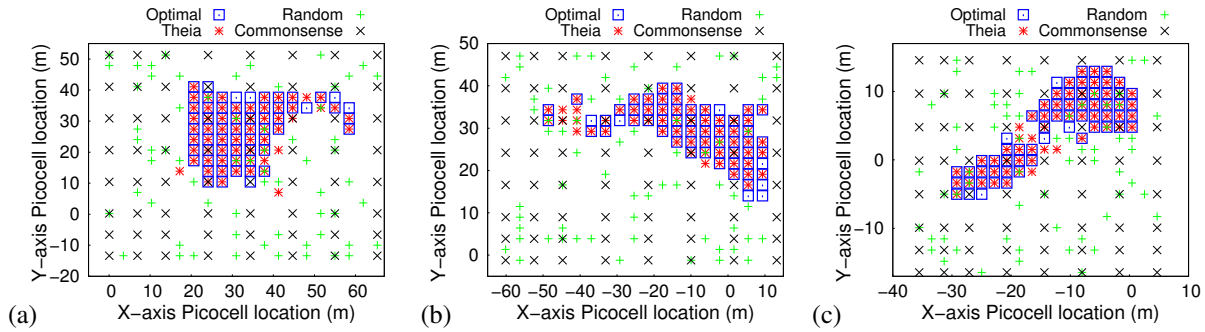

Figure 13: Picocell locations estimated by Random, Commonsense, Optimal, and *Theia* deployment strategies using the link-outage method.

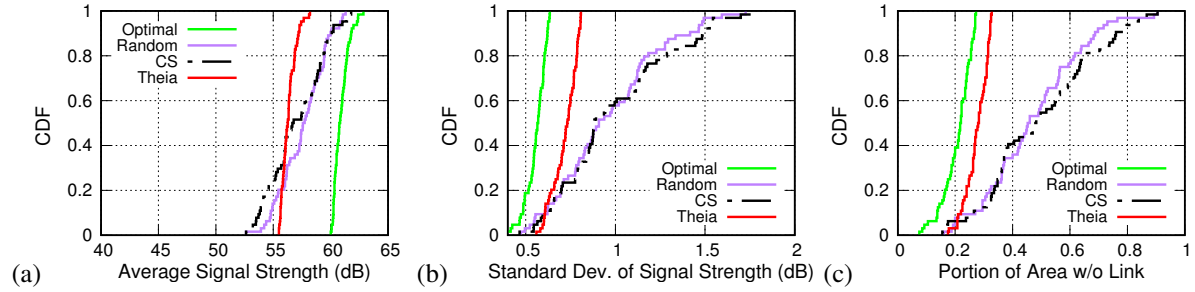

Figure 14: Deployment performance metrics for the environment A considering; (a) Average signal strength at a location. (b) Signal strength variation in the environment. (c) Distribution of portion of the area without any link.

frequency networks, but extending such models to mmWave frequencies has been challenging [50].

To address this issue, Argus [5] has developed a machine learning augmented system that accurately predicts mmWave signal reflection in less time and is generalizable across multiple indoor environments. However, it does not account for outdoor obstacles such as cars, buildings, and pedestrians that affect diffraction nor does it consider the large scale of outdoor environments. Existing works, such as Lumos 5G [51], have attempted to predict 5G outdoor signal using Google Maps and RNN, but these methods lack consistency in client features and do not generalize well across different outdoor environments. In contrast to these existing works, *Theia* addresses the challenge of obtaining accurate SRPs in outdoor environments by measuring SRPs from sparse locations in the environment using a drone-based platform. Furthermore, *Theia*'s deep learning model is trained using on the dataset collected by the drone, and predicts SRP across multiple outdoor environments accurately.

MmWave Network Deployment: Recall that 5G planning has been done using empirical techniques and thorough surveying by hand. However, it is very meticulous, expensive, and time-consuming. Also, in large-scale outdoor environments, it may not even be feasible to survey every nook and cranny by hand. But, these techniques are still being used by many phone companies and the Federal Communications Commission. One approach to reduce time and cost uses ray propagation simulators to determine how signals are affected by obstacles and give a detailed map of signal strengths in the environment [52], [53], [54]. However, these methods are limited by the specific frequencies simulated since NLOS signal reflectivity is frequency-dependent. Other methods like [5], [55], [56], [57] use reflection information collected in indoor environments at various frequencies to determine the best locations for network infrastructure. However, outdoor environments have many factors such as area-span, foliage, number and height of buildings, and pedestrian and automotive traffic, which indoor environments do not. By using *Theia*, data can be collected easily in varied outdoor environments then use its picocell deployment method to enable the placement of 5G infrastructure in an expedited manner without compromising performance.

VII. Conclusion

In this work, we present *Theia*, a system for accurate predictions of mmWave signal reflection profiles in outdoor environments from unmeasured vantage points by using a deep learning model trained on data collected from a drone-based survey platform. Our experimental results show that *Theia* performs well across multiple large-scale outdoor environments and can fine-tune to new environments with as little as 5 minutes of survey data. Theia's SRP prediction can accurately estimate picocell placements with optimal coverage and enable reliable and robust outdoor mmWave networks.

VIII. Acknowledgments

We sincerely thank the reviewers for their comments. This work is partially supported by the NSF under grants CAREER-2144505, CNS-1910853, and MRI-2018966.

REFERENCES

[1] J.P. Morgan, "The Future of 5G Adoption," 2021. [Online]. Available: https://www.jpmorgan.com/insights/research/future-of-5g-adoption

[2] Federal Communications Commision, "FCC Promotes Higher Frequency Spectrum for Future Wireless Technology," 2015. [Online]. Available: https://apps.fcc.gov

[3] Rangan, Sundeep, et al., "Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges," *Proceedings of the IEEE*, 2014.

[4] ExterNetworks, "Wireless Site Survey - Outdoor Survey for upto 100,000 Sq. Ft," 2020. [Online]. Available: https://www.extnoc.com/store/wireless-outdoor-survey-100000sqft/

[5] Hem Regmi, et al., "Argus: Predictable Millimeter-Wave Picocells with Vision and Learning Augmentation," in *ACM SIGMETRICS*, 2022.

[6] Timothy Hooks, et al., "VisualMM: Visual Data & Learning Aided 5G Picocell Placement," in *ACM HotMobile*, 2021.

[7] DJI, "Matrice 100 Quadcopter for Developers," 2015. [Online]. Available: https://www.dji.com/matrice100

[8] ——, "Guidance," 2015. [Online]. Available: https://www.dji.com/guidance

[9] Texas Instruments, "IWR1443 single-chip 76-GHz to 81-GHz mmWave sensor evaluation module," 2020. [Online]. Available: https://www.ti.com/tool/IWR1443BOOST

[10] IEEE P802.11 - Task Group ay, "Status of Project IEEE 802.11ay," 2020. [Online]. Available: https://www.ieee802.org/11/Reports/tgay_update.htm

[11] Mehrdad Soumekh, *Synthetic Aperture Radar Signal Processing*. John Wiley and Sons., Inc., 1999, vol. 7.

[12] Zhou Wang, et al., "Image Quality Assessment: From Error Visibility to Structural Similarity," *Trans. Img. Proc.*, 2004.

[13] Chao Dong, et al., "Image Super-Resolution Using Deep Convolutional Networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.

[14] Wenming Yang, et al., "Deep learning for single image super-resolution: A brief review," *IEEE Transactions on Multimedia*, 2019.

[15] Kevin de Haan, et al., "Deep-Learning-Based Image Reconstruction and Enhancement in Optical Microscopy," *Proceedings of the IEEE*, 2020.

[16] Saiprasad Ravishankar, et al., "Image Reconstruction: From Sparsity to Data-Adaptive Methods and Machine Learning," *Proceedings of the IEEE*, 2020.

[17] Florian Knoll, et al., "Deep-Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A Survey of the Current Approaches, Trends, and Issues," *IEEE Audio and Electroacoustics Newsletter*, 2020.

[18] e. a. Sarabandi, K., "Compact beam scanning 240GHz radar for navigation and collision avoidance," in *Micro- and Nanotechnology Sensors, Systems, and Applications III*, 2011.

[19] Mohammad Mosalanejad, et al., "Multilayer Compact Grid Antenna Array for 79 GHz Automotive Radar Applications," *IEEE Antennas and Wireless Propagation Letters*, vol. 17, no. 9, 2018.

[20] He Huang, et al., "5G MIMO Antenna Based on Vector Synthetic Mechanism," *IEEE Antennas and Wireless Propagation Letters*, vol. 17, no. 6, 2018.

[21] Sharnil Pandya, et al., "Analysis of NOMA-OFDM 5G wireless system using deep neural network," *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 2021.

[22] Tian-Hao Li, et al., "Learning the Wireless V2I Channels Using Deep Neural Networks," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, 2019.

[23] Nagarathna Ravi, et al., "Secure Deep Neural (SeDeN) Framework for 5G Wireless Networks," in *2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, 2019.

[24] Jiuxiang Gu, et al., "Recent advances in convolutional neural networks," *Pattern Recognition*, 2018.

[25] Chen-Lin Zhang, et al., "In Defense of Fully Connected Layers in Visual Representation Transfer," in *Advances in Multimedia Information Processing - Pacific-Rim Conference on Multimedia*, 2017.

[26] Keras, "Keras Applications," 2022. [Online]. Available: https://keras.io/api/applications/

[28] Mingxing Tan, et al., "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in *Proceedings of the 36th International Conference on Machine Learning*, 2019.

[27] Karen Simonyan, et al., "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[29] Mark Sandler, et al., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[30] Andrew G. Howard, et al., "Searching for MobileNetV3," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.

[31] Barret Zoph, et al., "Neural Architecture Search with Reinforcement Learning," *CoRR*, 2016.

[32] Mingxing Tan, et al., "MnasNet: Platform-Aware Neural Architecture Search for Mobile," *CoRR*, 2018.

[33] Jia Deng, et al., "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009.

[34] Chang-Fa Yang, et al., "A ray tracing method for modeling indoor wave propagation and penetration," in *IEEE Antennas and Propagation Society International Symposium. 1996 Digest*, 1996.

[35] DJI-SDK, "DJI Onboard SDK," 2022. [Online]. Available: https://github.com/dji-sdk/Onboard-SDK

[36] ——, "DJI Guidance SDK," 2022. [Online]. Available: https://github.com/dji-sdk/Guidance-SDK

[37] Texas Instruments, "DCA1000EVM Real-time data-capture adapter for radar sensing evaluation module," 2020. [Online]. Available: https://www.ti.com/tool/DCA1000EVM

[38] e. a. David M. Sheen, "Three-Dimensional Millimeter-Wave Imaging for Concealed Weapon Detection," *IEEE Transactions on Microwave Theory and Techniques*, vol. 49, no. 9, 2001.

[39] Matthew Lloyd, "UAV Forecast," 2022. [Online]. Available: https://www.uavforecast.com/

[40] PyTorch, "Torch.Optim," 2022. [Online]. Available: https://pytorch.org/docs/stable/optim.html#algorithms

[41] Open-Source, "PyTorch," 2022. [Online]. Available: https://pytorch.org/

[42] NVIDIA, "NVIDIA RTX A6000 GRAPHICS CARD," 2022. [Online]. Available: https://www.nvidia.com/en-us/design-visualization/rtx-a6000/

[43] Emmanouil Alimpertis, et al., "City-Wide Signal Strength Maps: Prediction with Random Forests," in *The World Wide Web Conference*, 2019.

[44] Theodore Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Prentice Hall PTR, 2001.

[45] Ahmed Iyanda Sulyman, et al., "Radio propagation path loss models for 5G cellular networks in the 28 GHZ and 38 GHZ millimeter-wave bands," *IEEE Communications Magazine*, vol. 52, no. 9, 2014.

[46] Remcom, "Wireless InSite: 3D Wireless Prediction Software," 2022. [Online]. Available: https://www.remcom.com/wireless-insite-em-propagation-software

[47] Siradel, "Software for Wireless Network and Smart City Planning," 2022. [Online]. Available: https://www.siradel.com/smart-city-planning/

[48] Altair Engineering, Inc., "Altair Feko Applications," 2022. [Online]. Available: https://www.altair.com/feko-applications/

[49] SonicWall, "WiFi Planner: Elevate Your WiFi User Experience with the Right Design," 2022. [Online]. Available: https://www.sonicwall.com/products/secure-wireless/wifi-planner/

[50] Xing, Yunchou, et al., "Indoor Wireless Channel Properties at Millimeter Wave and Sub-Terahertz Frequencies," in *2019 IEEE Global Communications Conference (GLOBECOM)*, 2019.

[51] Arvind Narayanan, et al., "Lumos5G: Mapping and Predicting Commercial MmWave 5G Throughput," in *ACM Internet Measurement Conference*, 2020.

[52] Ferdous Hossain, et al., "Indoor Millimeter-Wave Propagation Prediction by Measurement and Ray Tracing Simulation at 38 GHz," *Symmetry*, 2018.

[53] Remcom, "5G Simulation & MIMO Simulation Software," 2023. [Online]. Available: https://www.remcom.com/5g-mimo

[54] Siradel, "Smart City Planning ," 2023. [Online]. Available: https://www.siradel.com/smart-city-planning/

[55] Hem Regmi, et al., "Towards Deep Learning Augmented Robust D-Band Millimeter-Wave Picocell Deployment," *ACM SIGMETRICS Performance Evaluation Review*, vol. 50, no. 4, 2023.

[56] Sanjib Sur, et al., "Poster: Scoping Environment to Assist 60 GHz Link Deployment," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, 2015.

[57] Teng Wei, et al., "Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors," in *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, 2017.