IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received December 7, 2020, accepted December 28, 2020, date of publication December 31, 2020, date of current version January 7, 2021.

Digital Object Identifier 10.1109/ACCESS.2020.3048583

FadeNet: Deep Learning-Based mm-Wave Large-Scale Channel Fading Prediction and Its Applications

VISHNU V. RATNAM[1], (Member, IEEE), HAO CHEN[1], (Member, IEEE), SAMEER PAWAR[1], BINGWEN ZHANG[1], CHARLIE JIANZHONG ZHANG[1], (Fellow, IEEE), YOUNG-JIN KIM[2], SOONYOUNG LEE[2], MINSUNG CHO[2], AND SUNG-ROK YOON[2]

[1]Standards and Mobility Innovation Laboratory, Samsung Research America, Plano, TX 75023, USA
[2]Network Division, Network Automation Group, Samsung Electronics Company Ltd., Suwon 16677, South Korea

Corresponding author: Vishnu V. Ratnam (vishnu.r@samsung.com)

ABSTRACT Accurate prediction of the large-scale channel fading is fundamental to planning and optimization in 5G millimeter-wave cellular networks. The current prediction methods, which are either too computationally expensive or inaccurate, are unsuitable for city-scale cell planning and optimization. This paper presents FadeNet, a convolutional neural-network enabled alternative for predicting large-scale fading with high computation speed and accuracy. By using carefully designed input features and neural-network architecture to capture topographical information, FadeNet accurately predicts the large-scale fading from a base station to each location in its coverage area. Evaluations on realistic data, derived from millimeter-wave cells across multiple cities, suggest that FadeNet can achieve a prediction accuracy of 5.6 decibels in root mean square error. In addition, by leveraging the parallel processing capabilities of a graphics processing unit, FadeNet can reduce the prediction time by $40X - 1000X$ in comparison to industry prevalent methods like ray-tracing. Generalizations of FadeNet, that can handle variable topographies and base station heights, and its use for optimal cell site selection are also explored.

INDEX TERMS Cell planning, channel modeling, convolutional networks, deep learning, large scale fading, mm-Wave, pathloss, U-net.

I. INTRODUCTION

The wireless signal strength from a base-station (BS) experiences significant variation as the receiver (RX) moves, due to interactions with scattering objects in the propagation environment. This effect, known as fading, is usually modeled as the combination of a large scale fading (LSF) component,[1] which varies slowly in time, frequency and space, and a small scale fading component, which accounts for fast random variations in time, frequency and space. Accurate modeling and prediction of the LSF component, in particular, is of vital importance for planning and optimization in cellular networks. With an accurate LSF model, operators can identify coverage holes, find the best locations to install new BSs, and optimize deployed BS configuration parameters (e.g., antenna heights, tilts, transmit power etc.) – a task referred to as network planning and optimization [1]. This is especially crucial for the imminent planned deployment for 5-th generation (5G) millimeter-wave (mm-Wave) networks where the use of higher carrier frequencies leads to a high dynamic range in LSF values.

Conventionally, there are two types of approaches to model the LSF in wireless channels, namely, statistical and deterministic. The first approach involves modeling the LSF statistically, with examples including the Okumura-Hata model [2], COST 231 model [3], [4], WINNER II model [5], QuadRiGa [6], 3GPP spatial channel model [7], etc. In this approach, the LSF from a transmitter (TX) to a RX is typically calculated by a distance-based propagation equation plus a random shadowing value. The randomness is aimed at achieving convergence in distribution of the LSF values rather than exactly predicting the LSF value for a given RX. Thus, while the cost and computational complexity of this approach are negligible, the resulting LSF models are not very accurate

The associate editor coordinating the review of this manuscript and approving it for publication was Cesar Vargas-Rosales.

[1]The LSF component here characterizes impact of both the distance-dependent pathloss and shadow fading. Some papers also refer to this combined term as pathloss.

since the models do not take into account the surrounding radio frequency (RF) environment between the TX and RX.[2] At mm-Wave frequencies, the LSF is deeply affected by the surrounding RF environment, including buildings, trees, terrain etc. Neglecting these effects, the LSF estimation errors from statistical channel models can be as large as 20 decibels (dB) or more. The second approach for modeling the LSF is to use deterministic methods like ray-tracing [8], [9], uniform theory of diffraction [10], Maxwell's equations [11], finite difference time domain method [12] etc. In this approach, detailed maps of the surrounding RF environment are maintained and used for predicting the LSF at each RX deterministically. In ray-tracing, for example, a large number of out-going rays are simulated from the TX, allowing them to be reflected, scattered and diffracted by the surrounding environment. Subsequently, for each RX point, all the rays that reach its vicinity are identified. Correspondingly, the details of Azimuth angle Of Arrival (AOA), Azimuth angle Of Departure (AOD), Zenith angle Of Arrival (ZOA), Zenith angle Of Departure (ZOD), power and delay for each ray/path between the TX and RX can be captured. Although the deterministic methods can accurately predict the LSF values, the cost of the required detailed RF maps, the computation time and the memory requirements of these methods can be prohibitively large. For example, the required three-dimensional (3-D) maps for ray-tracing may cost around $\$50/km^2$ to purchase and computation of LSF values for a 500m $\times$ 500m area using full 3-D ray-tracing may take up to 10 minutes. These large overheads make this approach less attractive for the network planning and optimization of a large city. Consequently several techniques have been proposed to bring down this computational complexity via dimensionality reduction [13], space partitioning [14], [15], using coarse simulation with interpolation [16], [17], constraining ray interactions and power levels [18], adopting parallel processing methods [19] etc. A detailed survey of such ray-tracing enhancements can be found in [20], which can bring down the LSF computation time by up to a factor of 10.

As an exciting alternate approach, this paper studies the use of supervised deep learning (DL) based techniques for modeling the LSF in a BS's coverage area. With deep neural networks being universal function approximators, this approach shows potential for achieving good modeling accuracy. Furthermore, deep neural networks are a natural fit for parallelization on graphics processing units (GPUs), thus promising high computation speeds. Finally, unlike the deterministic methods where several system parameters require manual tuning, DL based approaches can train directly on the data, thus reducing manual effort in system calibration. The following section summarizes the current progress in the literature on such DL techniques and our contributions.

[2]While the COST231 model does use the height information of buildings between the TX and RX, it still is quite inaccurate in practice.

A. RELATED WORKS AND OUR CONTRIBUTIONS

For a single TX and RX (point-to-point) framework, a significant amount of work has been done on supervised learning based LSF or pathloss[3] prediction. In [21], a comprehensive overview of DL techniques for several communication problems is presented. In [22], a general framework for training neural networks for pathloss prediction is proposed, albeit without considering the surrounding RF environment information. In [23], the authors use image segmentation to classify satellite images of an area of interest as urban, semi-urban, rural, free space and forest, and then use the corresponding COST-231 statistical models to calculate the point-to-point pathloss value. In [24], a Multi-layer Preceptron (MLP) is used to predict the point-to-point pathloss using as inputs the TX-RX distance and local parameters for terrain and vegetation, such as the terrain clearance angle and vegetation type surrounding the RX. In [25], a 1-layer MLP is proposed to predict the error between the statistical channel models and the measured pathloss, using as inputs the predicted pathloss and the distance between the TX and RX. In [26], support vector regression was used to predict the point-to-point pathloss value using the TX-RX distance, terrain elevation, horizontal, vertical angles and the prediction from Okumura-Hata model as inputs. In [27], an MLP is trained to predict the point-to-point LSF value using the building and terrain features as inputs. In [28], an MLP is designed to predict the point-to-point LSF using the TX-RX distance, TX/RX heights, diffraction loss for the line-of-sight (LoS) path and crossed distance in each clutter type as inputs. In [29], the LSF for the air-to-air channel between unmanned aerial vehicles was predicted using TX/RX altitude, distance, LoS feature and elevation angle as inputs. In [30], the authors use deep convolutional neural networks to predict the point-to-point LSF using some environmental features surrounding the RX. In [31], an MLP is trained to predict the point-to-point LSF using as input a classification of the local topography into several clutter types. However, these prior works have several limitations. Firstly, these solutions do not sufficiently exploit the surrounding RF environment information. Secondly, the solutions are trained and tested on a very limited set of topographies. Finally, there has been no prior work on the modeling of LSF from one TX to many RX points over a large region of interest (RoI), such as the TX coverage area. Such an RoI based prediction is vital for many cellular applications, such as network coverage prediction, automatic site selection, BS parameter optimization etc.

In this paper, we propose a supervised learning based solution for one-shot prediction of the LSF from one BS to each grid point in a large RoI (see Fig. 1), by using the building, terrain and foliage height information for the RoI as input features. We treat each grid point of the RoI as a

[3]Depending on the reference, pathloss may either refer to the complete LSF or only the distance dependent component of LSF (not accounting for shadowing by buildings or objects). In this paper we shall use the latter definition.

pixel, and represent the input features as well as the LSF prediction for the RoI as images. We then harness the technological maturity of state-of-the-art neural network solutions for image processing tasks [32], to model the LSF prediction problem as an image regression problem. Since the number of grid points in the RoI can be very large and since the input features in the vicinity of a grid point will have a greater impact on its LSF, we adopt convolutional neural networks as an ideal candidate for the regression task. As shall be shown, our proposed solution reduces the LSF computation time by a massive **40X - 1000X**, in comparison to conventional deterministic methods. The contributions of the paper are as follows:

- We propose a novel deep-learning solution to predict the LSF from a BS to a large RoI in one shot, with high accuracy, low cost and very low computation overhead.
- In addition, we also propose a LoS-aided deep-learning solution for LSF prediction with improved performance.
- The proposed solutions use contextual information for the whole RoI, including the heights and locations of buildings, foliage and terrain, and are trained and tested on a massive data set obtained from 3 different cities in USA, covering many different topographies.
- We also propose several base-line prediction approaches and evaluate performance in realistic scenarios.
- We present extensions of FadeNet for dealing with variable TX antenna height and arbitrary geographical topographies, and exemplify the application of FadeNet for the optimal BS site selection problem.

The organization of the paper is as follows: the problem formulation and the neural network details such as input feature encoding, network architecture, output encoding, training methodology etc. are discussed in Section II; the baseline solutions for LSF prediction are detailed in Section III; the performance evaluation of the different LSF prediction solutions is presented in Section IV; extensions of FadeNet to realize additional functionalities are discussed in Section V; the practical use of LSF prediction for cell-site planning is discussed in Section VI; the limitations of FadeNet and future directions of research are discussed in Section VII; and finally the conclusions of the work are summarized in Section VIII.

Notation: In this paper, matrices are represented by bold upper-case letters, sets are represented as calligraphic letters and scalars are represented by light lower-case letters. In addition, $[\mathbf{A}]_{i,j}$ represents the (i,j)-th element of a matrix $\mathbf{A}$ and $|\mathcal{A}|$ represents the cardinality of a set $\mathcal{A}$.

II. FadeNet

We define the RoI for a BS as a square region of dimension 512 × 512m with the BS located at the center, such that it encapsulates the typical coverage area of the BS. However, the dimensions and the square shape are only representative, and the methodology in the paper can also be extended to arbitrary RoIs. To leverage image processing neural networks, the

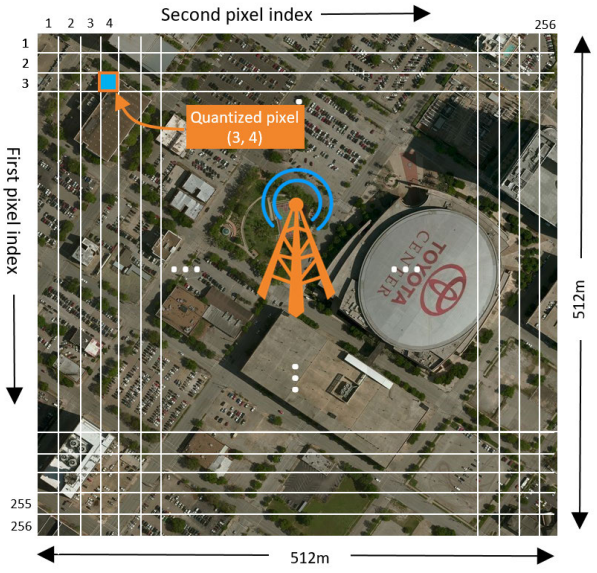

FIGURE 1. An illustration of the quantized RoI for a typical BS (Toyota Center, Houston, USA).

RoI is quantized into a grid of pixels of size 2 × 2 m, where each grid 'pixel' has a representative RX point at its center, as illustrated in Fig. 1. Thus, the RoI can be represented as a 256 × 256 matrix or image of RX points. Note that the BS is located at the bottom right corner of pixel (128,128). The BS antenna height is assumed to be $h_{\mathrm{BS}} = 7$m while each RX point is assumed to have an antenna height of $h_{\mathrm{RX}} = 1.5$m above the terrain. The extension to variable BS height shall be discussed later in Section V-B. The aim of FadeNet is to predict the LSF from the BS to all the representative RX points in the RoI, in one shot. In the following subsections, we discuss the input features, network architecture and the training methodology for FadeNet.

A. INPUT FEATURES FOR THE PURE DL SOLUTION

As mentioned previously, the RF environment strongly impacts the signal propagation at mm-Wave frequencies. Consequently, the environment information within an RoI, including terrain height, building height (above terrain level), and foliage height (above terrain level) are used as input features in FadeNet. Note that such environmental information is commercially available for most cities from map vendors such as Microsoft Bing, Google Maps etc. Each of the terrain, building, and foliage heights for an RoI r are encoded as images/matrices $\tilde{\mathbf{T}}_r, \tilde{\mathbf{B}}_r, \tilde{\mathbf{F}}_r$ of dimension 256 × 256, as illustrated in Fig. 2. Here the values at pixel (i,j), i.e., $[\tilde{\mathbf{T}}_r]_{i,j}, [\tilde{\mathbf{B}}_r]_{i,j}, [\tilde{\mathbf{F}}_r]_{i,j}$ are the maximum height (in meters) of the terrain, buildings and foliage, respectively, within the corresponding 2 × 2m quantized area of the RoI in Fig. 1. To equalize the dynamic ranges of these different inputs, they are further normalized as:

$$\mathbf{T}_r = \left(\tilde{\mathbf{T}}_r - [\tilde{\mathbf{T}}_r]_{128,128}\right)/25 \tag{1a}$$

$$\mathbf{B}_r = \tilde{\mathbf{B}}_r/100 \tag{1b}$$

$$\mathbf{F}_r = \tilde{\mathbf{F}}_r/10. \tag{1c}$$

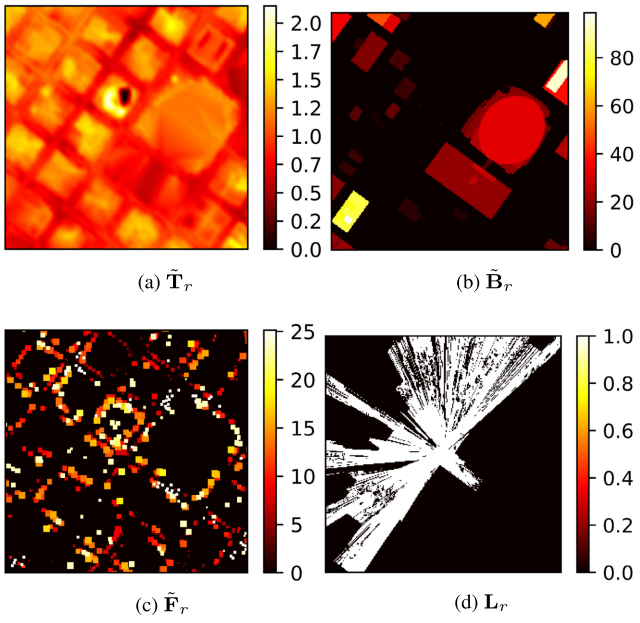

(a) $\tilde{\mathbf{T}}_r$ (b) $\tilde{\mathbf{B}}_r$

(c) $\tilde{\mathbf{F}}_r$ (d) $\mathbf{L}_r$

FIGURE 2. An illustration of the pre-normalized 256 × 256-pixel inputs $\tilde{\mathbf{T}}_r, \tilde{\mathbf{B}}_r, \tilde{\mathbf{F}}_r$ (in meters) and the additional LoS input $\mathbf{L}_r$ (refer to Section II-B) for the RoI in Fig. 1.

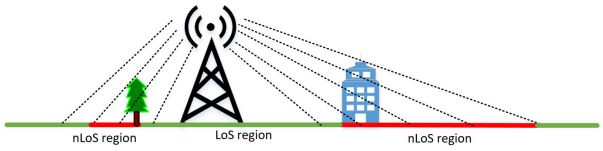

FIGURE 3. An illustration of the LoS and non-LoS (nLoS) regions for an example scenario.

Here the coefficient 25 in (1a) is the 95 percentile value of $[\tilde{\mathbf{T}}_r]_{i,j} - [\tilde{\mathbf{T}}_r]_{128,128}$, computed over the pixels (i, j) of many training RoIs r. Similarly, $\{100, 10\}$ are the 95 percentile values for $\tilde{\mathbf{B}}_r$ and $\tilde{\mathbf{F}}_r$, albeit when considering only pixels with non-zero values of $\tilde{\mathbf{B}}_r$ and $\tilde{\mathbf{F}}_r$, respectively. Note that adding a constant bias to the terrain height of all pixels does not affect the signal propagation or the LSF. Thus to avoid having many such equivalent inputs, the second term in (1a) is used to ensure that the terrain height at the BS, which is located near pixel (128,128), is always 0.

B. ADDITIONAL INPUT FOR THE LoS-AIDED DL SOLUTION

Mm-Wave signals suffer from strong attenuation losses by obstacles such as buildings, foliage, etc, that block the signal. Thus, the LSF at an RX point is highly correlated with the presence or absence of an unobstructed straight line path between the BS and RX point, called the line-of-sight (LoS) path. The presence/absence of the LoS path in an example scenario is illustrated in Fig. 3 for convenience. To improve the performance of FadeNet, we also explore the use of the LoS feature as input, in addition to the topographical features discussed in Section II-A. The LoS feature for an RoI r is encoded as a binary matrix/image $\mathbf{L}_r$, as exemplified in Fig. 2d, with $[\mathbf{L}_r]_{i,j} = 1$ if the corresponding 2 × 2m block has an LoS path to the BS and $[\mathbf{L}_r]_{i,j} = 0$ otherwise. Several algorithms exist for quickly computing $\mathbf{L}_r$ from the BS height, RX height, and the topographical inputs $\tilde{\mathbf{T}}_r, \tilde{\mathbf{B}}_r, \tilde{\mathbf{F}}_r$, such as the GIS 'viewshed' algorithm [33].

Although the LoS feature is a deterministic function of the other inputs $\tilde{\mathbf{T}}_r, \tilde{\mathbf{B}}_r, \tilde{\mathbf{F}}_r$, it has sharp variations and thus can be difficult to approximate for the hidden layers of a neural network. Thus, using it as an explicit additional input can potentially improve network performance, as shall be verified in Section IV. Note that LoS computation is also an important, albeit computationally inexpensive, sub-step of ray-tracing. Therefore, we shall refer to the FadeNet solution that only uses the inputs from Section II-A as the 'pure DL solution' while the one that additionally uses the LoS feature as the *LoS-aided DL solution*.

C. NEURAL NETWORK ARCHITECTURE

As mentioned in the introduction, we treat the LSF prediction for the RoI r as an image regression task, where the input images such as $\mathbf{B}_r, \mathbf{T}_r, \mathbf{F}_r$, etc, are fed as inputs to the neural network of FadeNet to generate the LSF prediction output image $\hat{\mathbf{Y}}_r$. As with the inputs, the desired format for the LSF prediction output $\hat{\mathbf{Y}}_r$ for an RoI r is also a matrix/image of dimension 256 × 256, with $[\hat{\mathbf{Y}}]_{i,j}$ being the predicted LSF (in dB) at the RX point in the center of (i, j)-th pixel. Since the dimensions of the inputs and the output are quite large (256 × 256) and since the output for a pixel is predominantly determined by input features in the neighboring pixels, we explore the use of a convolutional neural network for this regression task. We adopt a 28-layer convolutional network for FadeNet, as illustrated in Fig. 4, which is a variant of a popular image segmentation network called the U-Net [34]. For the uninitiated, a primer on the basic operations in a convolutional network can be found in [35], which are skipped here for brevity.

The output generation procedure for an RoI r using FadeNet is as follows. First, the input images for the RoI are stacked together to form a 256 × 256 × C tensor and are fed to layer 1 of the network, where C is the number of input images/channels ($C = 3$ for pure DL and $C = 4$ for the LoS-aided DL solution). The input to each layer then undergoes a two-dimensional convolution with the layer's filters, is post-processed and is fed as input to the subsequent layer. The proposed network has several strided convolution layers (layers 1-12), which extract the high-level features from the input tensor via dimensionality reduction. This is followed by several de-convolution layers (layers 13-28) which then construct the higher dimensional output image $\hat{\mathbf{Y}}_r$. There also exist skip connections between the convolution and de-convolution layers of the same dimensions, to pass the detailed local information from the inputs into the output construction phase. In Fig. 4, the network layers are shown as arrows with the color coding representing the convolution filter dimensions and the operations within each layer, and the dimensions of the inputs and outputs for these different layers

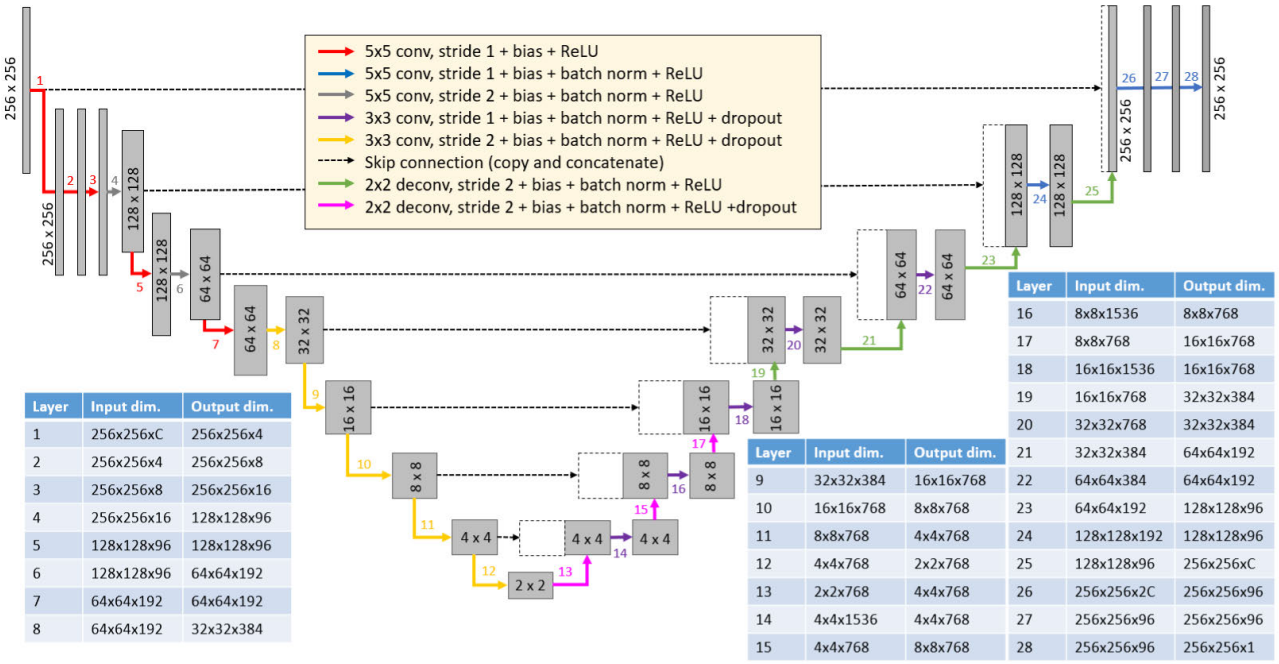

Layer	Input dim.	Output dim.
1	256x256xC	256x256x4
2	256x256x4	256x256x8
3	256x256x8	256x256x16
4	256x256x16	128x128x96
5	128x128x96	128x128x96
6	128x128x96	64x64x192
7	64x64x192	64x64x192
8	64x64x192	32x32x384

Layer	Input dim.	Output dim.
9	32x32x384	16x16x768
10	16x16x768	8x8x768
11	8x8x768	4x4x768
12	4x4x768	2x2x768
13	2x2x768	4x4x768
14	4x4x1536	4x4x768
15	4x4x768	8x8x768

Layer	Input dim.	Output dim.
16	8x8x1536	8x8x768
17	8x8x768	16x16x768
18	16x16x1536	16x16x768
19	16x16x768	32x32x384
20	32x32x768	32x32x384
21	32x32x384	64x64x192
22	64x64x384	64x64x192
23	64x64x192	128x128x96
24	128x128x192	128x128x96
25	128x128x96	256x256xC
26	256x256x2C	256x256x96
27	256x256x96	256x256x96
28	256x256x96	256x256x1

FIGURE 4. An illustration of the network architecture adopted for FadeNet. The network has 28 layers, represented by colored arrows described in the legend, and the corresponding layer inputs and outputs are represented as rectangles. The dimensions of these inputs and outputs are tabulated below the network structure. In the legend, '$a \times a$ conv, stride b' and '$a \times a$ deconv, stride b' represent 2D-convolution and transpose 2D-convolution with an $a \times a$ filter, respectively, followed by an output decimation by a factor of b. Similarly, 'ReLU' is the rectified linear unit activation function, and 'dropout', 'batch norm' are commonly used regularization techniques [36], [37]. We use a dropout probability of 0.8 for the layers $8-12$ and 0.4 for layers ≥ 13, and C in the Table refers to the number of input features/images per RoI.

are also summarized in the tables. Despite the similarity to U-Net, there are also several modifications in the proposed architecture owing to differences in the associated learning tasks. Firstly, since FadeNet is designed for an image regression task, unlike U-Net which is designed for an image segmentation task, we add several additional stride-1 convolution layers at the output (layers 26-28). Secondly, unlike U-Net, FadeNet uses multiple input images and the ground-truth output is significantly different from the inputs in appearance. Consequently, we also add additional stride-1 convolution layers at the input (layers 1-3) to enable mixing of the inputs before dimensionality reduction, as illustrated in Fig. 4. The network dimensions and parameters chosen for FadeNet have been engineered by an extensive and intricate parameter and network tuning process. We shall henceforth use the variable $\mathbf{W}$ to represent the trainable parameters (weights, biases etc.) of the neural network architecture.

D. GROUND TRUTH AND NETWORK OUTPUT

Note that the FadeNet output can be expressed as: $\hat{\mathbf{Y}}_r = f(\text{inputs}, \mathbf{W})$, where the inputs are from Section II-A and/or II-B, and $\mathbf{W}$ is a matrix of trainable weights associated with the neural network. Supervised deep-learning solutions typically require a large, labelled training data-set, and consequently the *ground-truth* LSF values for a large number of RoIs are required for training the FadeNet weights $\mathbf{W}$. Since obtaining actual LSF values via field test measurements for many RoIs is practically infeasible, we use ray-tracing to generate the ground-truth LSF values for the training RoIs. Note that, while being computationally expensive, ray-tracing can provide an accurate prediction of the true LSF and is thus a good substitute ground truth [38]–[40]. Furthermore, since FadeNet only requires the ray-tracing results during the initial training phase, the cost and computational benefits of FadeNet over conventional ray-tracing are not precluded. The ground-truth for RoI r is also encoded as a 256×256 dimension image $\mathbf{Y}_r$, with $[\mathbf{Y}]_{i,j}$ being the ray-tracing LSF (in dB) at the RX point in the center of the (i, j)-th pixel. Such a matrix/image can be easily computed from the prediction result of any commercial raytracing tool for the whole RoI r. An example of the ground truth LSF $\mathbf{Y}_r$ and prediction output $\hat{\mathbf{Y}}_r$ for the RoI in Fig. 1 are illustrated in Fig. 5 for convenience. Note that the ray-tracing prediction may not be available for all pixels of the RoI, such as pixels inside buildings or pixels to which signal propagation is difficult. Such regions are represented with black color in Fig. 5a. Because of this limitation, FadeNet never encounters labeled training data for pixels inside buildings and can give inaccurate results for such pixels. Therefore, we shall intentionally post-process the FadeNet output $\hat{\mathbf{Y}}_r$ to black out the pixels with $[\mathbf{B}_r]_{i,j} > 0$, as illustrated in Fig. 5b.[4]

[4]This is a limitation of the ray-tracing tool used rather than a limitation of FadeNet. By using an alternate ground-truth tool that can generate indoor LSF values, FadeNet can be trained to predict indoor LSF values as well.

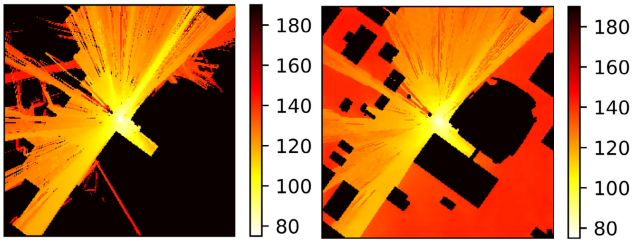

(a) $\mathbf{Y}_r$ (in dB) (b) $\hat{\mathbf{Y}}_r$ (in dB)

FIGURE 5. An illustration of the 256 × 256-pixel ray-tracing based ground truth $\mathbf{Y}_r$ and the predicted LSF output $\hat{\mathbf{Y}}_r$ for the RoI in Fig. 1.

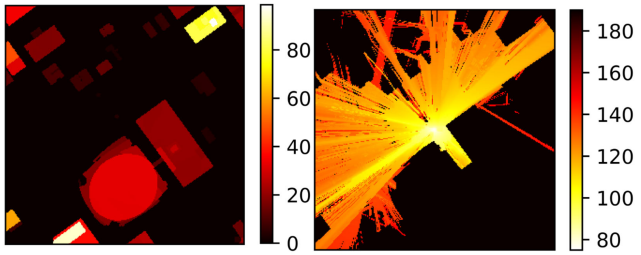

(a) $\tilde{\mathbf{B}}_{\bar{r}}$ (b) $\mathbf{Y}_{\bar{r}}$

FIGURE 6. An illustration of the building height input $\tilde{\mathbf{B}}_{\bar{r}}$ and the ray-tracing based ground truth $\mathbf{Y}_{\bar{r}}$ for the augmented RoI $\bar{r}$, obtained by a horizontal flip followed by a 270° clockwise rotation of the RoI in Fig. 1.

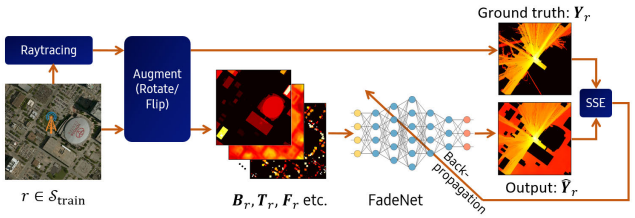

FIGURE 7. An illustration of the training procedure for FadeNet for a sample RoI.

Since the ground-truth is generated from ray-tracing, FadeNet can be interpreted as a cost and computationally efficient approximator to ray-tracing. However, FadeNet also paves the way for *transfer learning* [41], a technique wherein a neural network trained on ray-tracing can be fine-tuned using a limited amount of field test based actual LSF measurements. While beyond the scope of this paper, such a transfer-learned FadeNet can potentially out-perform the accuracy of ray-tracing based LSF prediction. To quantify the mismatch between the prediction output $\hat{\mathbf{Y}}_r$ and the ground truth $\mathbf{Y}_r$, we use the root mean square error (RMSE) as the performance measure. For an RoI r, the RMSE can be computed as:

$$\text{RMSE}(r) = \sqrt{\sum_{(i,j)\in\mathcal{V}_r} \left([\hat{\mathbf{Y}}_r]_{i,j} - [\mathbf{Y}_r]_{i,j}\right)^2 \Big/ |\mathcal{V}_r|}, \qquad (2)$$

where $\mathcal{V}_r$ is defined as the subset of pixels in the RoI for which valid ray-tracing based ground-truth exists, i.e., the non-dark pixels of Fig. 5a.

E. DATA AUGMENTATION AND TRAINING METHODOLOGY

It is well known that larger the training data set, the better is the performance and robustness of the resulting trained neural network. Thus to successfully train FadeNet, we collected 2500+ sample RoIs with the associated ground-truth LSF images from the city of Houston, USA, covering a wide variety of topographies. An augmentation method is employed to further increase the sample size, wherein new training RoIs are generated from existing ones via a combination of horizontal flipping and/or clockwise rotation (by 90°, 180° or 270°) of the input and ground truth images. One such example of horizontal flip + rotation by 270° of $\tilde{\mathbf{B}}_r$ in Fig. 2b and $\mathbf{Y}_r$ in Fig. 5a is illustrated in Fig. 6. By using this augmentation technique, the data samples are augmented by a factor of 8 to generate an aggregate of 20K+ sample RoIs. Apart from increasing the training data size, another benefit of the data augmentation method is that the resulting FadeNet is invariant to rotation and flipping of inputs, which is a desired characteristic. In other words, rotation and flipping operations can remove the bias in prediction results over different directions from a BS.

Among the set of all RoI samples $\mathcal{S}_{\text{all}}$, a randomly selected 80% subset is chosen as the training set $\mathcal{S}_{\text{train}}$. The FadeNet network weights $\mathbf{W}$ are trained over $r \in \mathcal{S}_{\text{train}}$ to minimize the sum square error (SSE) between $\hat{\mathbf{Y}}_r$ and $\mathbf{Y}_r$, defined as:

$$\text{SSE} = \sum_{r\in\tilde{\mathcal{S}}_{\text{train}}} \text{RMSE}(r)^2 |\mathcal{V}_r|,$$

where $\tilde{\mathcal{S}}_{\text{train}}$ is the augmented superset of $\mathcal{S}_{\text{train}}$. Of the remaining 20% samples, half are chosen as the validation set $\mathcal{S}_{\text{val}}$ and are used for selecting the best network parameters and for early stopping the training process [42]. The final 10% samples (called the testing set $\mathcal{S}_{\text{test}}$) are then used to evaluate the network performance in Section IV. The training methodology for FadeNet is depicted pictorially in Fig. 7.

III. BASELINE CLASSICAL LEARNING SOLUTION

As mentioned in Section I, there is very limited work on learning/AI-based, one-shot LSF prediction for multiple receive points in an RoI. Consequently, for quantifying the benefits of using DL for the LSF prediction problem, in this section we also formulate two baseline classical learning solutions, namely, (i) Distance-dependent prediction and (ii) Conditional least-squares prediction.

A. DISTANCE-DEPENDENT PREDICTION

This approach is based on the extensively used log-distance pathloss model [4], where the LSF (in dB) for a pixel (i, j) of any RoI is predicted by a distance-dependent formula:

$$[\hat{\mathbf{Y}}^{\text{dd}}]_{i,j} = \alpha + 10\eta \log_{10}\left[d_{i,j}\right], \qquad (3)$$

where α, η are tunable parameters representing the bias and pathloss exponent, and $d_{i,j}$ is the LoS distance between the BS, located at the bottom-right corner of pixel (128, 128), and

RX point. Thus, $d_{i,j}$ is computed as:

$$d_{i,j}^2 = 4(i-128.5)^2 + 4(j-128.5)^2 + (h_{\mathrm{BS}} - h_{\mathrm{RX}})^2,$$

where $h_{\mathrm{BS}}, h_{\mathrm{RX}}$ are the heights of the BS and RX (in meters). The optimal tunable parameters α^*, η^* are then obtained as the minimizers of the sum square prediction error over the augmented training set as:

$$\mathrm{argmin}_{\alpha,\eta}\Big\{ \sum_{r\in\tilde{\mathcal{S}}_{\mathrm{train}}} \sum_{(i,j)\in\mathcal{V}_r} \big([\hat{\mathbf{Y}}_r^{\mathrm{dd}}]_{i,j} - [\mathbf{Y}_r]_{i,j}\big)^2 \Big\}.$$

Note that this model only uses the distance from the BS as an input, and is thus reminiscent of the statistical prediction models discussed in Section I.

B. CONDITIONAL LEAST-SQUARES PREDICTION

As the distance from the BS $d_{i,j}$ is a fixed value for a pixel (i,j), the solution in Section III-A always generates the same LSF prediction $\hat{\mathbf{Y}}^{\mathrm{dd}}$ for any RoI. Such a topology-agnostic prediction is clearly sub-optimal. Among the topographical features, the presence or absence of the LoS path has a strong correlation with the LSF for a RX point, as discussed in Section II-B. Thus we also propose a conditional least-squares prediction solution, that uses the LoS feature $\mathbf{L}_r$ to predict LSF as:

$$[\hat{\mathbf{Y}}_r^{\mathrm{cls}}]_{i,j} = \begin{cases} [\mathbf{Z}^{(0)}]_{i,j} & \text{if } [\mathbf{L}_r]_{i,j} = 0 \\ [\mathbf{Z}^{(1)}]_{i,j} & \text{if } [\mathbf{L}_r]_{i,j} = 1, \end{cases} \tag{4}$$

where $\mathbf{Z}^{(0)}$ and $\mathbf{Z}^{(1)}$ are two tunable 256×256 matrices, representing the LSF values for pixels of the RoI under the nLoS and LoS scenarios, respectively. Note that unlike in (3), we now have two candidate values $[\mathbf{Z}^{(0)}]_{i,j}, [\mathbf{Z}^{(1)}]_{i,j}$ for each pixel (i,j) based on the LoS condition. Furthermore, unlike the distance dependent equation in (3) that only has two parameters α, η, here these matrix entries are more flexible and are computed as the least-square solutions of:

$$[\mathbf{Z}^{(0)}]_{i,j} = \mathrm{argmin}_a \Big\{ \sum_{r\in\tilde{\mathcal{S}}_{\mathrm{train}}|(i,j)\in\mathcal{V}_r} \big(a - [\mathbf{Y}_r]_{i,j}\big)^2 \big(1 - [\mathbf{L}_r]_{i,j}\big) \Big\}$$

$$[\mathbf{Z}^{(1)}]_{i,j} = \mathrm{argmin}_a \Big\{ \sum_{r\in\tilde{\mathcal{S}}_{\mathrm{train}}|(i,j)\in\mathcal{V}_r} \big(a - [\mathbf{Y}_r]_{i,j}\big)^2 [\mathbf{L}_r]_{i,j} \Big\}.$$

Note that in the first equation, a is the least-squares fit to the LSF values observed for pixel (i,j) over all training samples with $[\mathbf{L}_r]_{i,j} = 0$. Similarly in the second equation, a is the least-squares fit to the LSF values observed for pixel (i,j) over all training samples with $[\mathbf{L}_r]_{i,j} = 1$.

IV. EVALUATION

In this section, we evaluate and compare the performance of FadeNet: (i) without LoS input (pure DL solution) and (ii) with LoS input (LoS-aided DL solution) and the baseline solutions (iii) Distance dependent prediction and (iv) Conditional least-squares prediction. For ease of notation, the networks in (i) and (ii) shall be denoted as $\overline{\mathrm{FadeNet}_{\mathrm{H}}}$ and $\mathrm{FadeNet}_{\mathrm{H}}$, respectively. For the training and evaluation, 2500 RoIs are picked in the city of Houston, USA and the corresponding topographical maps, including building, terrain and foliage information are purchased from a map vendor. For generating the ground truth LSF values of these RoIs, we use a state-of-the-art commercial 3-D ray-tracing tool for the 28GHz frequency band that uses a ray-launching based technique [43]. To improve computation efficiency, we perform the ray-launching with a ray angular resolution of 0.1 degrees in both azimuth and elevation, and with the per ray upper bounds on the number of reflections, number of diffractions and maximum ray LSF of 4, 1 and 200dB, respectively. Furthermore, since we are only interested in LSF, we use omni-directional antennas at the BS and RX with antenna heights of $h_{\mathrm{BS}} = 7$m and $h_{\mathrm{RX}} = 1.5$m for the ray-tracing.

Of the 2500 labeled RoIs, FadeNet and the baseline solutions are trained over a set of $|\mathcal{S}_{\mathrm{train}}| = 2000$ RoIs. Of the remaining RoIs, 250 are used for validation of FadeNet, while a final 250 RoIs are used for comparison of the solutions' performance. While the training is performed with data augmentation, the validation and testing are performed without it. For quantifying the performance, we use the average RMSE as the metric, obtained as:

$$\mathrm{RMSE} = \sqrt{\frac{\sum_{r\in\mathcal{S}_*} \mathrm{RMSE}(r)^2 |\mathcal{V}_r|}{\sum_{r\in\mathcal{S}_*} |\mathcal{V}_r|}},$$

where $* =$ train, val or test, depending on the situation. The training and testing are performed on a Linux machine with 1TB RAM and a Tesla P40 Nvidia GPU. For solutions (ii) and (iv), the LoS input is generated using the GIS viewshed algorithm [33], which requires a computation time of 3 secs/RoI.[5]

For training FadeNet, we use the Adam optimizer with Xavier weight initialization and zero bias initialization. The learning rate is decayed from 0.001 to 0.00001 during the training process, and we adopt mini-batch training with a batch size of 8. The maximum number of epochs is set to 1000, but we use the early stopping criterion (based on lowest validation set RMSE) to prevent over-fitting in the training phase. The training and validation RMSE of $\mathrm{FadeNet}_{\mathrm{H}}$ along with the early stopping point are illustrated in Fig. 8a. For the baseline solutions, the optimal parameter values for the distance dependent LSF formula in (3) computed over $\tilde{\mathcal{S}}_{\mathrm{train}}$ are: $\alpha^* = 48.36$dB and $\eta^* = 3.21$. For these parameter values, the predicted LSF for any RoI is depicted in Fig. 8b. Similarly for the conditional least-squares solution, the optimal choice of the matrices $\mathbf{Z}^{(0)}, \mathbf{Z}^{(1)}$ after training on $\tilde{\mathcal{S}}_{\mathrm{train}}$ are depicted in Fig. 8.

The average RMSEs of these solutions on the training set and test set are depicted in Table 1. As evident from the

[5]While not explored here, a significant boost in compute speed may be possible by parallelizing the LoS computation for each pixel.

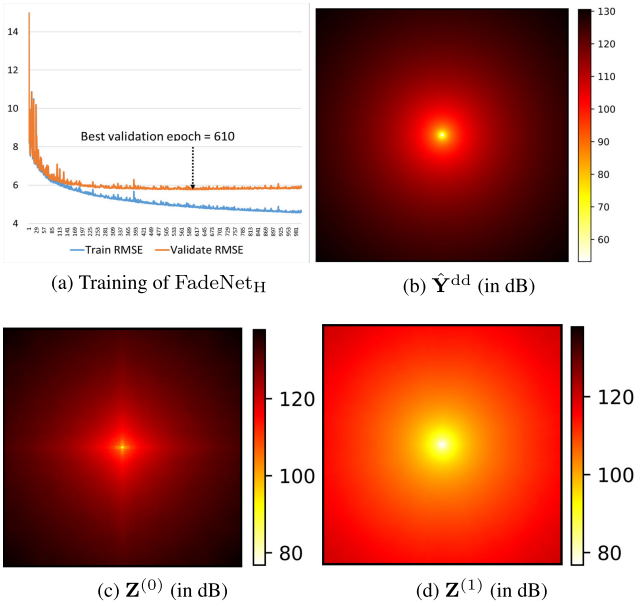

(a) Training of $\text{FadeNet}_{\text{H}}$ (b) $\hat{\mathbf{Y}}^{\text{dd}}$ (in dB)

(c) $\mathbf{Z}^{(0)}$ (in dB) (d) $\mathbf{Z}^{(1)}$ (in dB)

FIGURE 8. Training of the proposed solutions on the augmented training set $\tilde{\mathcal{S}}_{\text{train}}$. (a) depicts the training and validation RMSE of FadeNet during the training process. (b) depicts the 256 × 256 LSF prediction from Section III-A with optimal α^*, η^*. (c) and (d) depict the optimal conditional LSF predictions from Section III-B for the nLoS and LoS scenarios, respectively.

TABLE 1. Average test RMSE and the computation speed/test RoI of different LSF prediction solutions on Houston test set.

	Distance Dependant	Conditional Least-Squares	$\overline{\text{FadeNet}}_{\text{H}}$	$\text{FadeNet}_{\text{H}}$
Train RMSE	12.8 dB	8.6 dB	5.7 dB	**4.7 dB**
Test RMSE	12.8 dB	8.6 dB	7.1 dB	**5.8 dB**
Predict speed	0 s/RoI	3 s/RoI	$<$ 0.1 s/RoI	3 s/RoI

results, the test RMSE of the distance dependent prediction is quite high since the topographical information of the RoI is not exploited. Since the conditional least-squares prediction can exploit some topographical information encoded in the LoS input $\mathbf{L}_r$, it achieves a lower test RMSE. Note that the test RMSEs for these baseline solutions is still quite high and is identical to their train RMSEs, suggesting that these are high bias, low variance estimators [44]. From the result of the pure DL solution $\overline{\text{FadeNet}}_{\text{H}}$, we observe that FadeNet can out-perform the base-line solutions by $>$ 1.5 dB even without explicitly using the LoS information. However, the use of the LoS information as an additional input can boost the performance of FadeNet by a further 1.3 dB, as observed from the test RMSE of the LoS-aided DL solution $\text{FadeNet}_{\text{H}}$. While we observe some gap between the training RMSE and test RMSE of FadeNet, attempts to improve performance by reducing this gap (via regularization) were unsuccessful. The distribution of RMSE across different test RoIs for the

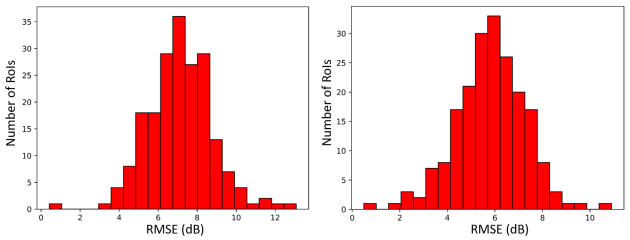

(a) $\overline{\text{FadeNet}}_{\text{H}}$ (b) $\text{FadeNet}_{\text{H}}$

FIGURE 9. Histogram of the RMSE per RoI RMSE(r) (in dB) for the test set $\mathcal{S}_{\text{test}}$, for each of the two FadeNet solutions.

two FadeNet solutions is depicted in Fig. 9. As can be seen from the results, the RMSE distribution is concentrated and is bounded below 12dB for the whole test set.

In addition to the RMSE, the compute time per RoI for each of these solutions is also depicted in Table 1, where we account for input generation time and the time for the forward pass through network but not the time required to load the network and data into the GPU. As observed, the pure DL FadeNet solution requires a compute time of only 0.1s/RoI while the LoS-aided solution requires 3s/RoI, due to computation of LoS feature. In comparison, the 3-D raytracing tool used for generating the ground truth values requires $>$ 120s to generate the LSF prediction for an RoI of this dimension. The authors have verified that several other popular commercial ray-tracing tools such as Wireless Insite, Atoll, etc., also require a similar compute time. We thus conclude that our DL solutions $\overline{\text{FadeNet}}$ and FadeNet reduce the LSF computation time by a massive **1000X** and **40X**, respectively, in comparison to conventional deterministic methods. This large speed boost facilitates several of new applications, as shall be discussed later in Section VI.

V. EXTENSIONS

In this section, we explore extensions of FadeNet to more generalized scenarios, including joint multi-city LSF prediction, accommodating for variable BS antenna height etc. Due to the superior performance of the LoS-aided DL solution FadeNet over the pure DL solution $\overline{\text{FadeNet}}$, here we shall restrict the comparisons to the former solution.

A. MULTI-CITY TRAINING VS. INDIVIDUAL TRAINING

In this subsection, we explore the application of FadeNet to multi-city LSF prediction. For this purpose, we collect 2500+ RoI samples each from 3 cities of USA: Houston, Sacramento and San Francisco. We use 250, 250 sample RoIs from each city for validation and testing, respectively, while the rest of the samples are used as the training set. Let us denote $\text{FadeNet}_{\text{H}}$, $\text{FadeNet}_{\text{S}}$ and $\text{FadeNet}_{\text{SF}}$ as the FadeNet networks trained exclusively on Houston, Sacramento and San Francisco data, respectively. In addition, we also define $\text{FadeNet}_{\text{H,S}}$ and $\text{FadeNet}_{\text{H,S,SF}}$ as the FadeNet networks that are trained on the joint training sets of Houston+Sacramento and all three cities, respectively. The performance of these

TABLE 2. Average RMSE of the different FadeNet networks, on the test sets of three cities.

Network	Houston	Sacramento	San Francisco
$\text{FadeNet}_{\text{H}}$	5.8 dB	5.8 dB	8.2 dB
$\text{FadeNet}_{\text{S}}$	5.9 dB	5.7 dB	8.4 dB
$\text{FadeNet}_{\text{SF}}$	6.3 dB	6.2 dB	7.2 dB
$\text{FadeNet}_{\text{H,S}}$	**5.6 dB**	**5.5 dB**	9.1 dB
$\text{FadeNet}_{\text{H,S,SF}}$	5.7 dB	5.6 dB	**7.0 dB**

trained networks on the test sets of the three cities are tabulated in Table 2. As observed from Table 2, the performance of $\text{FadeNet}_{\text{H}}$ is similar to that of $\text{FadeNet}_{\text{S}}$ on the Sacramento test set, but its performance is significantly worse than $\text{FadeNet}_{\text{SF}}$ on the San Francisco test set and vice versa. This is due to the similar terrain topography of Houston and Sacramento, which is significantly different from that of San Francisco. Furthermore, a larger fraction of buildings are sky-rises in San Francisco than in Houston or Sacramento. This suggests that a network trained on one city only generalizes well to cities with similar topographical features. We also observe that the network trained on both Houston and Sacramento $\text{FadeNet}_{\text{H,S}}$ yields slightly better test performance than $\text{FadeNet}_{\text{H}}$ and $\text{FadeNet}_{\text{S}}$ on the Houston and Sacramento test sets, respectively. This is due to the similarities of the two cities data, and the larger training set used by $\text{FadeNet}_{\text{H,S}}$.[6] Finally, we observe that $\text{FadeNet}_{\text{H,S,SF}}$ has better performance in comparison to $\text{FadeNet}_{\text{H}}$, $\text{FadeNet}_{\text{S}}$ and $\text{FadeNet}_{\text{SF}}$ on the Houston, Sacramento and San Francisco test sets, respectively. This suggests the our proposed FadeNet structure has sufficient capacity to learn intrinsic features of multiple cities. Thus if having separate networks for each type of city is infeasible, a good alternative is to jointly train FadeNet on all the cities data.

B. VARIABLE BS HEIGHT

The results in Section IV are applicable for the scenario where the BS antenna height is fixed ($h_{\text{BS}} = 7$m). Therefore, the BS antenna height is not fed as an input to the neural network. However, in a real deployment, the height of a mm-Wave BS may vary based on the available installation locations. In fact, field trials have shown that in real deployment scenarios, the BS antenna height values can vary from 2m to 50m. In this section, we therefore explore the generalization of FadeNet to the scenario where the BS height is variable. For this, we collect 3000 new training RoIs for the city of Houston, of which we have 600 RoIs uniformly distributed in each of the following BS height ranges: $2-5m, 5-10m, 10-20m, 20-30m, 30-50m$. The train, validation and test sets are chosen of sizes

[6]Note that a 2X increase in the training set size (for $\text{FadeNet}_{\text{H,S}}$) only improves performance by 0.2dB in comparison to $\text{FadeNet}_{\text{H}}$, $\text{FadeNet}_{\text{S}}$. Thus, due to the diminishing returns, we have restricted the sample size per city $\mathcal{S}_{\text{all}}$ to 2500.

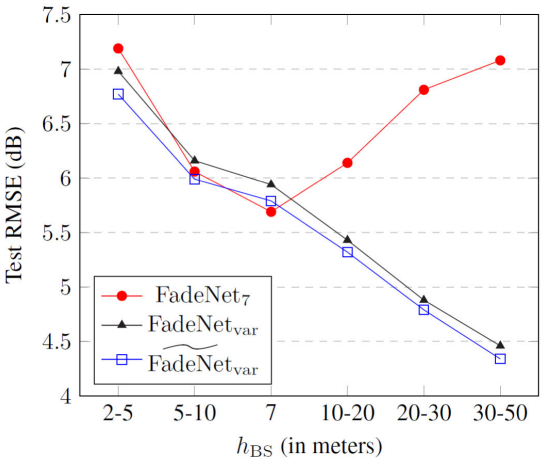

FIGURE 10. RMSE versus BS antenna height range for different FadeNet solutions.

2000, 300 and 700, respectively, while ensuring equal representation of each height range in these sets. Let us denote FadeNet_7 and $\text{FadeNet}_{\text{var}}$ as the FadeNets trained on the $h_{\text{BS}} = 7$m RoIs (from Section IV) and the new variable BS height RoIs, respectively. In addition we also introduce a new network $\widetilde{\text{FadeNet}}_{\text{var}}$ that, in addition to the inputs $\mathbf{B}_r, \mathbf{T}_r, \mathbf{F}_r, \mathbf{L}_r$ for an RoI r, also uses an antenna height image/matrix $\mathbf{H}_r$ as an input, and is trained on the new variable BS height RoIs. Here $\mathbf{H}_r$ is a 256×256 constant matrix, with each element being equal to the BS antenna height (in meters), normalized by 25.[7] Note that the network architecture for $\widetilde{\text{FadeNet}}_{\text{var}}$ is identical to Section II-C, with the number of inputs adjusted from $C = 4$ to $C = 5$.

The comparison of the performance of these networks on the variable height test set is illustrated in Fig. 10, for different height ranges. In addition, we also include the performance on the 7m Houston test set from Section IV. As observed from the results, FadeNet_7 performs well for 7m height range but has poor performance for other heights since it is trained only on $h_{\text{BS}} = 7$m data. In comparison, we see that $\text{FadeNet}_{\text{var}}$ has reasonable performance at all heights since it is trained on the variable height data, albeit without the antenna height as input. For a network that properly exploits BS antenna height as input, we expect its test RMSE to be lower than that of $\text{FadeNet}_{\text{var}}$ at all antenna heights and only slightly higher than $\widetilde{\text{FadeNet}}_7$ on the $h_{\text{BS}} = 7m$ test data. As evident from Fig. 10, $\widetilde{\text{FadeNet}}_{\text{var}}$ outperforms $\text{FadeNet}_{\text{var}}$ at all heights by 0.2 dB and is only 0.1 dB worse than FadeNet_7 at 7m height range, thus demonstrating its efficacy in dealing with an arbitrary BS height. As a side note, the lower test RMSE of $\widetilde{\text{FadeNet}}_{\text{var}}$ at larger BS heights is because the LoS coverage area within an RoI increases with BS height, and the LSF prediction for LoS locations is less erroneous.

[7]Several other encodings of the antenna height were also explored, but did not yield any performance improvement.

VI. USE CASE: OPTIMAL CELL SITE SELECTION

Many wireless operators are currently deploying network infrastructure to provide 5G mm-Wave cellular coverage. As discussed previously, the channel propagation at mm-Wave frequencies is strongly impacted by the surrounding RF environment. Therefore, selection of good site locations for deployment of mm-Wave BSs is of paramount importance to improve the network coverage while keeping the deployment cost low. One good approach for such selection in a deployment region is to: (i) determine a large set of candidate sites, (ii) predict the received signal power from each of the sites to the their surrounding area and (iii) determining a subset of sites that jointly maximize the signal coverage within the deployment region. Here we define coverage as the set of locations with received signal power above a certain threshold from any one of the active BSs. Note that the average received signal power from an omni-directional transmitter can be obtained from the LSF value as: $P_{\text{rx,dBm}} = P_{\text{tx,dBm}} - \text{LSF}$, where $P_{\text{tx,dBm}}$ is the transmit power. Thus the solutions discussed in this paper can be used for received power prediction in step (ii). Of the solutions discussed in the paper, statistical prediction models yield poor prediction results and the deterministic methods like ray-tracing are computation and time intensive. For example, for an area with 100 candidate sites, the time required for generating ray-tracing based LSF predictions is $\approx$ 4 hours. In comparison, the FadeNet solutions $\overline{\text{FadeNet}_\text{H}}$, FadeNet_H provide good accuracy and can perform the LSF prediction for the 100 candidate sites in 10s and 5 minutes, respectively, thus significantly reducing the turn-around time for network planning.

An illustrative demo of this use case using FadeNet is depicted in Fig. 11, where we have a 1km x 1km deployment area from Houston with 9 candidate site locations. FadeNet_H is used to predict the received signal power from each candidate site to the 512m x 512m area surrounding it. Every subset of these 9 candidate sites is a valid solution for deploying active BSs. For each solution of active BSs, we calculate the coverage as the fraction of pixels that have LSF smaller than 135 dB from at least one of the active BSs.[8] An exhaustive search is then performed over the BS subsets to find the smallest set of active sites that satisfies the 90% network coverage. For this toy demo, the optimal subset can be found in < 30 seconds.

VII. LIMITATIONS AND FUTURE DIRECTIONS

While FadeNet significantly improves the LSF prediction performance in comparison to the baseline approaches in the prior art, there are several opportunities to further improve the performance and address its drawbacks. These are discussed in this section and shall be explored in future work.

[8]The value of 135dB is just exemplary for this toy demo and is obtained using a back-of-the-envelope calculation with $P_{\text{tx,dBm}} = 30$dBm, a transmission bandwidth of 14MHz and minimum RX signal-to-noise-ratio of -6dB [45].

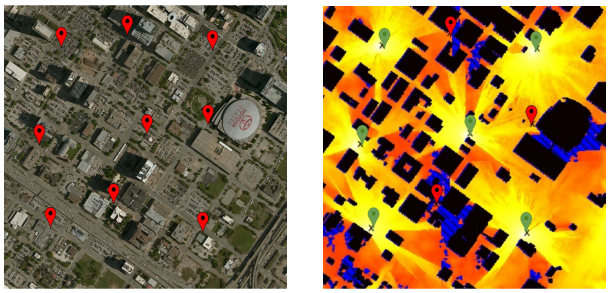

(a) 1×1 km deployment area (b) Predicted LSF and coverage

FIGURE 11. A demonstration of the optimal site selection to maximize the network coverage using FadeNet. (a) demonstrates the 1 × 1 km deployment area along with the candidate BS locations and (b) illustrates the selected 6 active sites with green markers. For (b) we also overlay a 500 × 500 pixel predicted LSF image, where the value at a pixel is the minimum LSF (in dB) from any of the 6 active sites. Pixels with no prediction or with best LSF > 135 dB are shaded with blue color (i.e., no coverage).

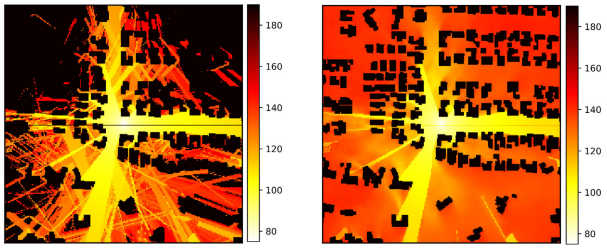

(a) Ground truth: $\mathbf{Y}_r$ (b) Prediction: $\hat{\mathbf{Y}}_r$

FIGURE 12. Comparison of $\mathbf{Y}_r$ and $\hat{\mathbf{Y}}_r$ (in dB) for a test RoI r for which FadeNet_H yields a large prediction error (RMSE = 8.2dB).

1) PREDICTING REFLECTIONS

It has been observed that FadeNet struggles to predict signal reflections properly, at least in some of the RoIs. This is evident from Fig. 12, where the ground-truth $\mathbf{Y}_r$ and the prediction result $\hat{\mathbf{Y}}_r$ (from FadeNet_H) are depicted for a high RMSE test RoI r. This behavior is somewhat expected since reflections by buildings are quite sensitive to the building surface, and their prediction using pixelated building height inputs can be quite difficult. Exploring mechanisms for better prediction of the signal reflections in FadeNet shall be explored in future work.

2) ACCOMMODATING FOR THE ANTENNA PATTERN

Since the focus of this paper is on LSF prediction, an omni-directional BS is considered and FadeNet returns, for each RX point in RoI, only the LSF values. Other channel parameters like angle-of-arrival (AoA) and angle-of-departure (AoD) of the dominant paths etc., which ray-tracing methods provide, are not predicted. Such parameters can be quite essential to predict how the received signal power in an RoI changes when the BS uses a directional antenna pattern. Due to the unavailability of these parameters, the use of FadeNet for scenarios with arbitrary BS antenna patterns is non-trivial and needs further investigation. Potential approaches for such a generalization include: using the antenna pattern as an input

to FadeNet, or predicting AoA, AoD as additional outputs with $\hat{\mathbf{Y}}_r$, and shall be explored in future work.

3) IMPROVING NETWORK ARCHITECTURE

Note that we model the LSF prediction as an image regression task, and propose a network architecture that is similar to the U-Net architecture [34] for this task. With the rapid advances in the image processing literature, there potentially exist other network architectures which may be better suited for the LSF prediction task, which can also be explored in future work.

VIII. CONCLUSION

This paper presents FadeNet, a deep learning based solution for predicting the LSF from a mm-Wave BS to the surrounding coverage area using the environment topographical features as inputs. From the results, we observe that FadeNet significantly outperforms statistical methods in terms of prediction accuracy, and reduces the computational complexity by $40-1000$X in comparison to deterministic methods like ray-tracing. We also conclude that the presence or absence of the line-of-sight path from the BS has a strong influence on the LSF value, and its use as an input feature to FadeNet can improve the prediction RMSE by 1.3 dB. Furthermore, FadeNet trained on cities with certain type of topographies only generalizes well to new cities with similar topographical features. We also conclude that LSF prediction is easier for line-of-sight points, and thus is easier for taller BSs which provide more line-of-sight coverage. Due to its prediction speed, FadeNet can also significantly reduce the turn around time of network planning tasks such as optimal cell site selection. Finally, we also conclude that there is scope for improving FadeNet to better predict the impact of signal reflections, and to account for variable BS antenna patterns.

REFERENCES

[1] P. V. Klaine, M. A. Imran, O. Onireti, and R. D. Souza, "A survey of machine learning techniques applied to self-organizing cellular networks," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 4, pp. 2392–2431, 4th Quart., 2017.

[2] M. Hata, "Empirical formula for propagation loss in land mobile radio services," *IEEE Trans. Veh. Technol.*, vol. VT-29, no. 3, pp. 317–325, Aug. 1980.

[3] E. Damosso and L. Correia. Digital Mobile Radio Towards Future Generation Systems: COST Action 231. European Commission. Brussels, Belgium. Accessed: 1999. [Online]. Available: https://books.google.com/books?id=setUHQAACAAJ

[4] A. F. Molisch, *Wireless Communications*, 2nd ed. Hoboken, NJ, USA: Wiley, 2010. [Online]. Available: https://www.ebook.de/de/product/12049967/molisch_wireless_communications_2e.html

[5] J. Meinil, P. Kysti, T. Jms, and L. Hentil, "WINNER II channel models," in *Radio Technologies and Concepts for IMT-Advanced*. Hoboken, NJ, USA: Wiley, 2010, pp. 39–92.

[6] S. Jaeckel, L. Raschkowski, K. Borner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3242–3256, Jun. 2014.

[7] *Study on Channel Model for Frequencies From 0.5 to 100 GHz*, document 3GPP TR 38.901, May 2017.

[8] M. C. Lawton and J. P. McGeehan, "The application of a deterministic ray launching algorithm for the prediction of radio channel characteristics in small-cell environments," *IEEE Trans. Veh. Technol.*, vol. 43, no. 4, pp. 955–969, Nov. 1994.

[9] S.-C. Kim, B. J. Guarino, T. M. Willis, V. Erceg, S. J. Fortune, R. A. Valenzuela, L. W. Thomas, J. Ling, and J. D. Moore, "Radio propagation measurements and prediction using three-dimensional ray tracing in urban environments at 908 MHz and 1.9 GHz," *IEEE Trans. Veh. Technol.*, vol. 48, no. 3, pp. 931–946, May 1999.

[10] W. M. O'Brien, E. M. Kenny, and P. J. Cullen, "An efficient implementation of a three-dimensional microcell propagation tool for indoor and outdoor urban environments," *IEEE Trans. Veh. Technol.*, vol. 49, no. 2, pp. 622–630, Mar. 2000.

[11] J.-M. Gorce, K. Jaffres-Runser, and G. de la Roche, "Deterministic approach for fast simulations of indoor radio wave propagation," *IEEE Trans. Antennas Propag.*, vol. 55, no. 3, pp. 938–948, Mar. 2007.

[12] J. W. Schuster and R. J. Luebbers, "Comparison of GTD and FDTD predictions for UHF radio wave propagation in a simple outdoor urban environment," in *IEEE Antennas Propag. Soc. Int. Symp. Dig.*, Jul. 1997, pp. 2022–2025.

[13] Z. Lai, H. Song, P. Wang, H. Mu, L. Wu, and J. Zhang, "Implementation and validation of a 2.5D intelligent ray launching algorithm for large urban scenarios," in *Proc. 6th Eur. Conf. Antennas Propag. (EUCAP)*, Mar. 2012, pp. 2396–2400.

[14] R. P. Kammaje and B. Mora, "A study of restricted BSP trees for ray tracing," in *Proc. IEEE Symp. Interact. Ray Tracing*, Sep. 2007, pp. 55–62.

[15] H. Mi, D. He, K. Guan, B. Ai, C. Liu, T. Shui, L. Zhu, and H. Mei, "Implementation and evaluation of ray-tracing acceleration methods in wireless communication," in *Proc. 14th Eur. Conf. Antennas Propag. (EuCAP)*, Mar. 2020, pp. 1–5.

[16] L. Azpilicueta, M. Rawat, K. Rawat, F. M. 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, May 2014.

[17] F. Casino, L. Azpilicueta, P. Lopez-Iturri, E. Aguirre, F. Falcone, and A. Solanas, "Optimized wireless channel characterization in large complex environments by hybrid ray launching-collaborative filtering approach," *IEEE Antennas Wireless Propag. Lett.*, vol. 16, pp. 780–783, 2017.

[18] M. Lecci, P. Testolina, M. Giordani, M. Polese, T. Ropitault, C. Gentile, N. Varshney, A. Bodi, and M. Zorzi, "Simplified ray tracing for the millimeter wave channel: A performance evaluation," 2020, *arXiv:2002.09179*. [Online]. Available: https://arxiv.org/abs/2002.09179

[19] J. Tan, Z. Su, and Y. Long, "A full 3-D GPU-based beam-tracing method for complex indoor environments propagation modeling," *IEEE Trans. Antennas Propag.*, vol. 63, no. 6, pp. 2705–2718, Jun. 2015.

[20] D. He, B. Ai, K. Guan, L. Wang, Z. Zhong, and T. Kurner, "The design and applications of high-performance ray-tracing simulation platform for 5G and beyond wireless communications: A tutorial," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 10–27, 1st Quart., 2019.

[21] S. M. Aldossari and K.-C. Chen, "Machine learning for wireless communication channel modeling: An overview," *Wireless Pers. Commun.*, vol. 106, no. 1, pp. 41–70, May 2019, doi: 10.1007/s11277-019-06275-4.

[22] Y. Zhang, J. Wen, G. Yang, Z. He, and J. Wang, "Path loss prediction based on machine learning: Principle, method, and data expansion," *Appl. Sci.*, vol. 9, no. 9, p. 1908, May 2019.

[23] S. Phaiboon, P. Phokharatkul, and P. Kittithamavongs, "Mobile path loss prediction with image segmentation and classification," in *Proc. Int. Conf. Microw. Millim. Wave Technol.*, Apr. 2007, pp. 1–4.

[24] E. Ostlin, H. Zepernick, and H. Suzuki, "Macrocell path-loss prediction using artificial neural networks," *IEEE Trans. Veh. Technol.*, vol. 59, no. 6, pp. 2735–2747, Jul. 2010.

[25] J. Isabona and V. M. Srivastava, "Hybrid neural network approach for predicting signal propagation loss in urban microcells," in *Proc. IEEE Region 10 Humanitarian Technol. Conf. (R HTC)*, Dec. 2016, pp. 1–5.

[26] R. Timoteo, D. Cunha, and G. Cavalcanti, "A proposal for path loss prediction in urban environments using support vector regression," in *Proc. Adv. Int. Conf. Telecommun. (AICT)*, Jan. 2014, pp. 119–124.

[27] S. P. Sotiroudis, K. Siakavara, and J. N. Sahalos, "A neural network approach to the prediction of the propagation path-loss for mobile communications systems in urban environments," *PIERS Online*, vol. 3, no. 8, pp. 1175–1179, 2007.

[28] M. Ayadi, A. Ben Zineb, and S. Tabbane, "A UHF path loss model using learning machine for heterogeneous networks," *IEEE Trans. Antennas Propag.*, vol. 65, no. 7, pp. 3675–3683, Jul. 2017.

[29] Y. Zhang, J. Wen, G. Yang, Z. He, and X. Luo, "Air-to-air path loss prediction based on machine learning methods in urban environments," *Wireless Commun. Mobile Comput.*, vol. 2018, pp. 8489326:1–8489326:9, Jun. 2018.
[30] N. Kuno and Y. Takatori, "Prediction method by deep-learning for path loss characteristics in an open-square environment," in *Proc. Int. Symp. Antennas Propag. (ISAP)*, Oct. 2018, pp. 1–2.
[31] L. Wu, D. He, B. Ai, J. Wang, H. Qi, K. Guan, and Z. Zhong, "Artificial neural network based path loss prediction for wireless communication network," *IEEE Access*, vol. 8, pp. 199523–199538, 2020.
[32] Y. Guo, Y. Liu, T. Georgiou, and M. S. Lew, "A review of semantic segmentation using deep neural networks," *Int. J. Multimedia Inf. Retr.*, vol. 7, no. 2, pp. 87–93, Nov. 2017.
[33] H. Haverkort, L. Toma, and Y. Zhuang, "Computing visibility on terrains in external memory," *ACM J. Experim. Algorithmics*, vol. 13, pp. 1–5, Feb. 2009, doi: 10.1145/1412228.1412233.
[34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham, Switzerland: Springer, 2015, pp. 234–241.
[35] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
[36] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014. [Online]. Available: http://jmlr.org/papers/v15/srivastava14a.html
[37] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn. (ICML)*, vol. 37, F. Bach and D. Blei, Eds. Lille, France: PMLR, Jul. 2015, pp. 448–456. [Online]. Available: http://proceedings.mlr.press/v37/ioffe15.html
[38] H. C. Nguyen, G. R. MacCartney, T. Thomas, T. S. Rappaport, B. Vejlgaard, and P. Mogensen, "Evaluation of empirical ray-tracing model for an urban outdoor scenario at 73 GHz E-band," in *Proc. IEEE 80th Veh. Technol. Conf. (VTC-Fall)*, Sep. 2014, pp. 1–6.
[39] A. Karstensen, W. Fan, F. Zhang, J. Ø. Nielsen, and G. F. Pedersen, "Analysis of simulated and measured indoor channels for mm-wave beamforming applications," *Int. J. Antennas Propag.*, vol. 2018, pp. 1–17, Jan. 2018, doi: 10.1155/2018/2642904.
[40] J.-Y. Lee, J.-H. Lee, and S.-C. Kim, "Improving the accuracy of millimeter-wave ray-tracing simulations by modeling roadside trees," *IEEE Antennas Wireless Propag. Lett.*, vol. 18, no. 1, pp. 162–166, Jan. 2019.
[41] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
[42] L. Prechelt, "Automatic early stopping using cross validation: Quantifying the criteria," *Neural Netw.*, vol. 11, no. 4, pp. 761–767, Jun. 1998. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0893608098000100
[43] G. Durgin, N. Patwari, and T. S. Rappaport, "An advanced 3D ray launching method for wireless propagation prediction," in *Proc. IEEE 47th Veh. Technol. Conf. Technol. Motion*, May 1997, pp. 785–789.
[44] S. Geman, E. Bienenstock, and R. Doursat, "Neural networks and the bias/variance dilemma," *Neural Comput.*, vol. 4, no. 1, pp. 1–58, Jan. 1992.
[45] Technical Specification Group Radio Access Network, *Study on New Radio Access Technology; Physical Layer Aspects*, document 3GPP TR 38.802, Tech. Rep. Release 14, Sep. 2017.

VISHNU V. RATNAM (Member, IEEE) received the B.Tech. degree (Hons.) in electronics and electrical communication engineering from IIT Kharagpur, Kharagpur, India, in 2012, where he graduated as the Salutatorian for the class of 2012, and the Ph.D. degree in electrical engineering from the University of Southern California, Los Angeles, CA, USA, in 2018. He currently works as a Senior Research Engineer with the Standards and Mobility Innovation Laboratory, Samsung Research America, Plano, TX, USA. His research interests include AI for wireless, mm-Wave and Terahertz communication, massive MIMO, and resource allocation problems in multi-antenna networks.

HAO CHEN (Member, IEEE) received the B.S. and M.S. degrees in information engineering from Xi'an Jiaotong University, Xi'an, in 2010 and 2013, respectively, and the Ph.D. degree in electrical engineering from the University of Kansas, Lawrence, KS, USA, in 2017. Since 2016, he has been a Research Engineer with the Standards and Mobility Innovation Laboratory, Samsung Research America. His research interests include network optimization, machine learning, and 5G cellular systems.

SAMEER PAWAR received the Ph.D. degree in electrical engineering and computer science from the University of California at Berkeley, in 2013. His research interests include machine learning and information theory. He is currently working as a ML Research Scientist at Facebook.

BINGWEN ZHANG received the B.E. degree from the University of Science and Technology of China, Hefei, China, in 2011, and the M.S. degree and the Ph.D. degree in electrical and computer engineering from Worcester Polytechnic Institute, MA, USA, in 2013 and 2017, respectively. His research interests include statistical learning and data mining.

CHARLIE JIANZHONG ZHANG (Fellow, IEEE) received the Ph.D. degree from the University of Wisconsin, Madison, WI, USA. He was with the Nokia Research Center, from 2001 to 2006, where he was involved in the IEEE 802.16e (WiMAX) standard and EDGE/CDMA receivers, and from 2006 to 2007, he was with Motorola, where he was involved in 3GPP HSPA standards. From 2009 to 2013, he served as the Vice Chairman of the 3GPP RAN1 Group and led the development of LTE and LTE-advanced technologies, such as 3-D channel modeling, UL-MIMO and CoMP, and carrier aggregation for TD-LTE. He is currently the Senior Vice President and Head of the Standards and Mobility Innovation Laboratory, Samsung Research America, where he leads research, prototyping, and standards for 5G cellular systems and future multimedia networks.

YOUNG-JIN KIM received the B.S. and Ph.D. degrees in electronic and electrical engineering from the Pohang University of Science and Technology (POSTECH), South Korea, in 2007 and 2013, respectively. He currently works as a Senior Research Engineer with the Network Division, Samsung Electronics, South Korea. His research interests include RAN optimization, admission control, and congestion management with machine learning for end-to-end service assurance in 5G network slicing.

SOONYOUNG LEE received the B.S. degree in electrical engineering and the M.S. and Ph.D. degrees in electrical engineering and computer science from Seoul National University, Seoul, South Korea, in 2005, 2007, and 2012, respectively. He is currently working with Samsung Electronics Company Ltd. His research interests include computer vision, data analysis, and AI.

MINSUNG CHO received the M.S. degree in electrical and telecommunication engineering from Sungkyunkwan University, South Korea, in 2000. He currently works as a Principal Research Engineer with the Network Division, Samsung Electronics Company Ltd. His research interests include RAN architecture, RAN network design and optimization, radio planning tool design, and air interface analysis.

SUNG-ROK YOON received the B.S. degree in electrical engineering and the M.S. and Ph.D. degrees in information and communications engineering from the Korea Advanced Institute of Science and Technology (KAIST), South Korea, in 2002, 2006, and 2011, respectively. He has been working with Samsung Electronics, South Korea, since 2011. His research interests include channel modeling, system KPI prediction, RAN optimization, RAN virtualization, and network automation of 5G cellular systems.