IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 17 April 2023, accepted 7 May 2023, date of publication 10 May 2023, date of current version 17 May 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3274810

RESEARCH ARTICLE

Deep Learning for Multi-User Proactive Beam Handoff: A 6G Application

FARIS B. MISMAR[1], (Senior Member, IEEE), ALPEREN GÜNDOĞAN[2], (Member, IEEE), ALIYE ÖZGE KAYA[3], (Senior Member, IEEE), AND OLEG CHISTYAKOV[4], (Member, IEEE)

[1]Nokia Bell Labs Consulting, Murray Hill, NJ 07974, USA
[2]Nokia Bell Labs, 81541 Munich, Germany
[3]Nokia Bell Labs, Murray Hill, NJ 07974, USA
[4]Nokia Bell Labs, 11742 Espoo, Finland

Corresponding author: Faris B. Mismar (faris.mismar@bell-labs-consulting.com)

ABSTRACT This paper demonstrates the use of deep learning and time series data generated from user equipment (UE) beam measurements and positions collected by the base station (BS) to enable handoffs between beams that belong to the same or different BSs. We propose the use of long short-term memory (LSTM) recurrent neural networks with three different approaches and vary the number of lookbacks of the beam measurements to study the performance of the prediction used for the proactive beam handoff. Simulations show that while UE positions can improve the prediction performance, it is only up to a certain point. At a sufficiently large number of lookbacks, the UE positions become irrelevant to the prediction accuracy since the LSTMs are able to learn the optimal beam based on implicitly defined positions from the time-defined trajectories.

INDEX TERMS Beamforming, deep learning, handoff, predictive, radio resource management, transfer learning.

I. INTRODUCTION

Since the introduction of the fifth generation of wireless networks (5G), user equipment (UEs) have become associated with serving beams instead of serving base stations. For stationary UEs, beams can improve the signal strength as measured by these UEs in contrast to a conventional BS with no beamforming. Beamforming in 5G allows the UE reference symbols to benefit from the beamforming gain—an opportunity the reference symbols missed in prior generations of wireless networks. However, several challenges emerge with mobility. These challenges can be summarized in either blockage or handoff interruption (or both). A proactive approach to beam assignment to moving UEs can help reduce the impact of these challenges and improve the reliability in these networks [1].

Handoffs are radio resource management (RRM) mobility procedures used to transfer the UE session from one base station (BS) to the other. In 5G, the concept of beam switching (or *handoff*) is introduced where a UE session is served by a different beam instead of its current one based on the radio measurements of the reported beams from UE. Beam switching can happen within the same BS (intra-BS) or between two different BSs (inter-BS). While it is straightforward for the UE to switch to a different beam in an intra-BS beam switching scenario, a random access procedure is often needed for the inter-BS beam switching scenario. In this case, the UE detects the most performance-optimal beam and sends back a physical random access mapped to the identifier of the beam. This procedure is not instantaneous since the UE has to wait for the random access procedure to take place before it can proceed [2]. During this time, further radio degradation in the current serving beam may happen and the UE may indicate a radio link failure (RLF) and perform a re-establishment procedure [3]. This procedure leads to longer service interruption as the UE now has to perform a random access procedure. It becomes an opportunity for an intelligent data-driven approach to proactively prepare the beam for every UE in transition. This is in order to avoid

The associate editor coordinating the review of this manuscript and approving it for publication was Di Zhang.

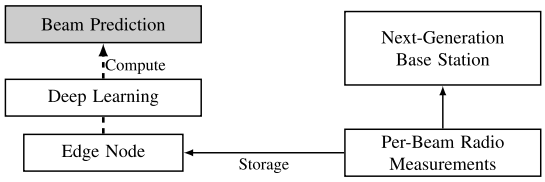

FIGURE 1. Next-generation base station uses an edge node for storage (solid) and compute (dashed) to proactively assign beams to UEs using radio measurements.

interruption in the flow of its data as a result of this beam switching.

With the development of technology beyond 5G, the introduction of artificial intelligence and machine learning (AI/ML) is inevitable: the data-driven approaches to solve problems that are mathematically intractable using AI/ML have witnessed established roots to industry standards. Specifically, there are study items on enhancements for data collection for 5G [4] and AI/ML air interface [5]. Deep learning, a specialized AI/ML technique, has shown viable reusability and reduction in computation cycles through "transfer learning" [6]. Transfer learning is a deep learning approach that uses a pre-trained model for one task as a starting point for another that performs a similar task, which can also be re-trained. It has been used with measurable success, as we also validate in this paper. Deep learning methods are a specialized form of machine learning that uses neural network architectures. Owed to their non-linear activation functions, these learning methods can learn arbitrarily complex relationships between inputs and outputs and have a certain level of robustness against noise. Also, due to the multiple layers of neural networks (or *depth*) and number of nodes in each layer (or *width*), they can autonomously extract learning features making the task of learning more time-efficient. Deep learning can perform time series classification. In this type of classification, training data is time-indexed and the learning must abide to the temporal relationship between observations in the data.

The use of deep learning in performing predictive handoff or beam switching serves other purposes with regards to the evolution towards the sixth generation of wireless networks (6G): on the one hand, it makes a solid use case for multi-access edge computing, where both data storage and compute power are brought near to the base station [7]. On the other hand, successful predictions can eliminate any potential disruption in the flow of data to the UE [8], [9], thus improving on the throughput and latency. These two measures remain quintessential performance measures as wireless networks evolve towards 6G. Fig. 1 shows how an edge compute and storage node can benefit the next-generation base stations in storing data and performing AI/ML-driven computations related to beam prediction and proactive beam handoffs.

A. RELATED WORK

A single user moving in a trajectory of two base stations (BSs) was studied in [1] and [10]. In [1], the objective was to minimize blocking probability using a gated recurrent unit (GRU) deep learner. We study multiple users moving in different trajectories with more than two BSs. Further, GRUs have the potential to overfit, a problem that can be avoided with models with more complexity, such as the long short-term memory deep neural network, which we propose. In [10], the assumption on the trajectory is that it is a fixed turning scene with a predetermined finite angle and radius. In our case, the trajectories are not predetermined and are based on two-dimensional streets with the motion of UEs being stochastic in both the direction and speed. Further, several papers successfully used these long short-term memory deep neural networks in areas not related to wireless networks such as [11], where different data-driven approaches were applied to proactively determine the useful life of Lithium-ion batteries and thus avoid probable damages.

The industry standards [3], [12] for mobility procedures can be divided as higher- and lower-layer mobility. Higher-layer mobility requires the Radio Resource Configuration (RRC) protocol to initiate handover procedure. The configuration to the UE can be provided earlier (i.e., conditional handover) or on time (i.e., baseline handover and dual active protocol stack handover) [3]. The handover in RRC layer is triggered if the beam received signal power of non-serving BS is higher than that of the beam of a serving BS by a certain offset and for a given duration [3]. In lower-layer mobility (LLM), the handover decision is performed by medium access control (MAC) layer of the air interface protocol stack based on the received beam measurements. MAC layer uses certain signaling information known as the "control elements," which are sent to the UE to trigger the serving beam change. UEs can be configured to report the beams of both serving BS and non-serving BS, and thus a beam change can be performed in both intra-BS or inter-BS scenarios. The configuration of the LLM mobility procedure is provided to the UE during the RRC procedure in advance as in conditional handover. In inter-BS mobility the decision to trigger LLM is made by the serving BS and conditions for this decision have not been defined in the standards. Thus, this provides room for novelty in terms of how this procedure can be implemented and we propose the predictive capability using prior measurements.

An alternative to beam switching was proposed in [8] through use of deep learning to perform band switching. In essence, channels of different frequency bands transmitted off the same BS (e.g., millimeter wave (mmWave) and sub-6 GHz) possess certain correlation characteristics that can be exploited to predict whether a UE switching from one frequency band to another would be successful. The use of UE coordinates was suggested as a means to improve the prediction performance. While we also propose the use of UE reported coordinates, we limit our dependence on the coordinates through the use of time series since time series captures the best serving beam of a given UE in its trajectory as a function of time. This implicitly captures the channel impact due to the location of the UE.

We note that in this paper, it is the 6G evolution of the base station to include compute capabilities that we exploit, and not any specific radio measurement changes, though these may prove helpful once the standards are finalized.

B. CONTRIBUTIONS

This paper makes the following specific contributions:

1) Demonstrate the importance of per-UE deep learners in BSs as opposed to one single learner per BS.
2) Show that knowledge of UE coordinates can improve prediction performance but up to a certain point in history.

II. SYSTEM MODEL

We consider a system composed of a multiple base stations (BS) with multiple transmit (i.e., downlink) and a single receive (i.e., uplink) antenna. Let there be M_T transmit antennas per BS. UEs are independently scattered in the service areas of these BSs and are in motion defined by a trajectory. Let the set $\mathcal{B}$ be defined as the set of BSs that cover the entire service area.

The received signal of the i-th UE on the b-th beam from the serving BS for a given subcarrier is therefore given by:

$$y_i = \mathbf{h}_{b,i}^* \mathbf{f}_{i,b}^\star x_i + \sum_{j \neq i} \mathbf{h}_{b,j}^* \mathbf{f}_{i,b}^\star x_j + n_i, \qquad i \in \{1, 2, \ldots, u\} \tag{1}$$

where x_i is the transmitted signal of the i-th UE normalized such that $P_{x_i} := \mathbb{E}|x_i|^2 = 1$ and $\mathbf{h}_{b,i} \in \mathbb{C}^M$ is the channel vector connecting the unique beam b to the i-th UE. We assume that the channel state information represented by these channel vectors $\mathbf{h}_{b,.}$ are perfectly known at the serving BS and UE, which is a common practice for transmissions at the mmWave frequency range where time division duplex (TDD) mode of operation is used and channel reciprocity is exploited. Further, we define these beams in a way that make them uniquely identifiable across all BSs. Let there be a total of $u > 0$ UEs in the system. The power-optimal beamforming vector $\mathbf{f}_{i,b}^\star$ is selected from a quantized analog beamforming codebook $\mathcal{F}$ with a finite cardinality:

$$\mathbf{f}_{i,b}^\star = \arg\max_{\mathbf{f} \in \mathcal{F}} |\mathbf{h}_{b,i}^* \mathbf{f}|^2. \tag{2}$$

The last two terms in (1) correspond to the inter-beam interference and additive white noise, which is sampled from a zero-mean Gaussian probability density function and a variance of σ^2. The received signal power can be computed as $P_{y_i} = |h_{b,i}^* f_{i,b}^\star|^2$, which when measured on the reference symbols on the synchronization channels per beam, is known as the "SS-RSRP" [13].

A. FADING

Since the UEs are in motion, they are expected to experience fading. In the presence of fading, we assume that the channel observes two types of fading: 1) large scale: shadow fading (i.e., log-Normal) and 2) small scale: Doppler fading depending on the movement speed of the UEs. The impact of the fading is captured in the channel state information.

B. BEAM SWITCHING

As mentioned in Section I, the change of the serving beam (i.e., the beam switching (or *handoff*) procedure) can be implemented at several layers in the air interface protocol stack. However, as in (1), the power-optimal beam is computed through a search over a codebook while the channel is assumed constant (i.e., within its coherence time). With beamforming, the channel coherence time increases due to directional reception [14]. The beam coherence time for a moving UE is approximately given by [8] and [15]:

$$t_{\text{coherence},i} \approx \frac{D_i}{v_i \sin \alpha_i} \frac{\Theta}{2}, \tag{3}$$

where D_i is the Euclidean distance of the i-th UE from the serving BS and Θ is the beamwidth of the serving beam measured in radians, v_i is the speed of the UE on the trajectory, and α_i is the angle between the direction of travel and the direction of the BS. UEs are expected to have different coherence times since they have different locations and distances from their respective serving BSs.

III. PROBLEM FORMULATION

The industry standards [12] provide a UE-specific identifier known as the BS "radio network temporary identifier" (C-RNTI). The C-RNTI is assigned to the UE during its random access procedure by the serving BS as shown in Fig. 2. It remains with the UE as long as it is served by that BS. This temporary identifier makes building a deep learning model per BS-UE pair possible. Further, standards [13] also provide identifiers per beam per BS. However, as there are multiple BSs, we propose one way to make the beam identifier unique across the network through a binary shift left and add operation, with the shift amount being equal to $\log_2 |\mathcal{B}|$.

This motivation allows us to formulate the problem of proactively selecting the optimal beam at the radio frame number n and the i-th UE:

$$b_i^\star[n+1] := \arg\max_{b \in \mathcal{B},\ \ell \in [0, |\mathcal{L}|-1]} \hat{M}_b(n; i, \ell) \tag{4}$$

where $\hat{M}_b(\cdot)$ is the deep learning estimate of the likelihood of the predicted beam to be for the optimal b-th beam using ℓ lookback beam measurements assigned to the i-th user. The lookback values selected from the ordered set $\mathcal{L}$. This formulation allows us to treat both intra- and inter-BS beam handoff alike.

From (4), we are training a multi-class classifier based on time series to predict the optimal beam for $(n+1)$, which means the beam for the *next* radio frame. While the offset term can be set to more than 1, we choose 1 which corresponds to the next radio frame. The main reason of predicting the optimal beam for the next radio frame only is that the beam switching (of the UE) to the target beam has to be completed within 1 ms for LLM according to the standards [16].

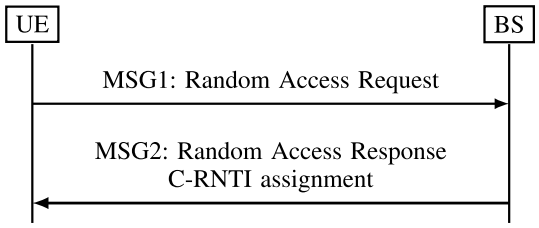

FIGURE 2. Assignment of the C-RNTI as part of the random access procedure defined in the standards.

IV. DEEP LEARNING FOR TIME SERIES

In this section we explain how time series data y_t can be used in deep learning. Particularly, we study the multi-variate time series case (i.e., where y_t is a function of many learning features measured at time t for a given UE).

A. TIME SERIES

Time series y_t is a sequence of data points indexed in *discrete* time t. This time can be in steps as small as milliseconds depending on the network configurations, which is the case when data points are collected from radio measurements. The v-th learning feature ($v \in \{0, 1, \ldots, N-1\}$) is denoted in a column vector format as: $\mathbf{x}^{(v)} := [x_t^{(v)}]_{t=0}^{M-1} \in \mathbb{R}^M$, where $x_t^{(v)}$ is the value of the v-th learning feature measured at the time instance t for the UE being served during this time instance.

In the context of our problem, y_t is the power-optimal beam identifier as measured at the radio frame t for a UE moving along an arbitrary trajectory. It is used to construct a column vector $\mathbf{y} := [y_t]_{t=0}^{M-1} \in \mathbb{R}^M$. All $\mathbf{x}^{(v)}$ and $\mathbf{y}$ are joined on the parameter t to construct the time-indexed dataset. At any given t and for a given UE, this dataset is a design matrix of joined $[\mathbf{x}^{(v)} \,|\, \mathbf{y}]$. For all UEs, a similar matrix can be constructed from this by stacking these design matrix instances based on time. This stacked matrix is referred to as $\mathbf{D} \in \mathbb{R}^{Mu \times n}$, where $u > 0$ is the number of UEs in the network as defined in (1). Given that the dataset contains a supervisory signal, $\mathbf{y}$, we may reconstruct this multi-variate time series problem as a supervised learning problem with the supervisory signal being a time-shifted version of $\mathbf{y}$. We discuss this as well as the impact of feature engineering further in this section.

Horizon is the number of time shifts in the *future* that we would like to predict. The horizon is a function of the learning features extending back in time, while it looks forward in time. In (4), our horizon is set to 1 as motivated earlier.

Embedding is transformation of a single dimension of temporal sequence to a multi-dimension space. By setting up a sequence of delays, we can treat each *past* value as an additional spatial dimension in the input space. If we look back in history by $\delta > 0$ time steps, then we have an embedding dimension of δ.

B. FEATURE ENGINEERING

There are three types of feature engineering that we consider: 1) shifting and lagging, 2) differencing, and 3) scaling.

1) SHIFTING AND LAGGING

To improve the predictability of a time series, a time shift of the dataset $\mathbf{D}$ is applied block-wise with past and future shift values. That is, the dataset becomes $[\mathbf{D}_{t-k} | \ldots | \mathbf{D}_t | \ldots | \mathbf{D}_{t+\ell}]$, where $k > 0$ is the lag and $\ell > 0$ is the lead time shift. This causes the dataset to have an additional total of $\ell + k$ column vectors, which represent the embedding dimension as defined earlier. This time shift operation also creates undefined values which are often dropped from the dataset causing a reduction in the number of *rows* by $\ell + k$.

2) DIFFERENCING

A stationary time series is a series the statistics of which do not depend on the time at which the series is observed. Differencing helps reduce any trend in the time series, and therefore introduces stationarity and improves predictability. Let us define the j-th order difference $d^{(j)}(t) := x_t - x_{t-j}, j \in \{1, 2, 3, \ldots M\}$ and apply it column-wise on $\mathbf{D}$. This leads to a new matrix $[\mathbf{\Delta}_{t-k} | \ldots | \mathbf{0} | \ldots | \mathbf{\Delta}_{t+\ell}]$, with $\ell+k+1$ columns. Combined with the shifting and lagging step, the dataset $\mathbf{D}$ thus has a total of $2(\ell + k + 1)$ column vectors, each of the dimension Mu for a shape $\mathbf{D} \in \mathbb{R}^{(M-(\ell+k)) \times 2(\ell+k+1)}$. We set $N := 2(\ell + k + 1)$ and $M' := Mu - (\ell + k)$ for simplification and thus $\mathbf{D} \in \mathbb{R}^{M' \times N}$ after the differencing operation.

3) SCALING

Scaling the learning features is done through the transformation function $\varphi : \mathbf{D} \mapsto \tilde{\mathbf{D}}$, where the scaling of values between [0, 1] is essential to avoid driving the neural network activation functions deep into saturation. We do not scale the supervisory signal vector.

C. SUPERVISED LEARNING AND PREDICTION

While the dataset $\mathbf{D}$ contains a supervisory signal $\mathbf{y}$ which is the time series at present time t, predictions are interested in *future* values of $\mathbf{y}$. To address this discrepancy, we define a lookahead $L \in \{0, 1, 2, \ldots\}$ where the supervisory signal becomes $\mathbf{y}_L := [y_{t+L}]_{t=0}^{M'-1} \in \mathbb{R}^{M'}$. This way, the supervised learning problem is established with respect to a lookahead value of the supervisory signal $\mathbf{y}_L$.

Combined with feature engineering, we can construct the time-series supervised learnable dataset $\mathbf{D} := [\tilde{\mathbf{D}} \,|\, \mathbf{y}_L]$. We use this definition of $\mathbf{D}$ as our learnable dataset moving forward, which is later reshaped into a tensor, to train the deep learning model as we see later in this section.

Prediction: If we denote the true supervisory signal at time t as $\mathbf{y}_L$, then the predicted signal is denoted as $\hat{\mathbf{y}}_L$. Since the supervisory signal $\mathbf{y}_L$ is categorical, the prediction problem becomes a minimization of the categorical cross-entropy of the probability density $p(\cdot)$ of $\hat{\mathbf{y}}_L$, which is defined as:

$$L(\mathbf{y}_L, \hat{\mathbf{y}}_L) := -\sum_{i=0}^{M'-1} \sum_{b \in \mathcal{F}} \mathbb{1}[[\mathbf{y}_L]_i = b] \cdot \log_2 p([\hat{\mathbf{y}}_L]_i = b). \tag{5}$$

D. DEEP LEARNING

Deep learning is a specific type of machine learning that uses several layers of connected computational elements called ''perceptrons,'' which are threshold functions with weights and are the basic unit of a neural network. These units can either be 1) connected where information flows forward between layers (i.e., from the input layer to the intermediate layer to to the output layer) or 2) recurrent where information can also flow through feedback connections. As any machine learning algorithm, deep learning also aims at minimizing a cost function. This is done through computing the gradient with respect to the weights of this function.

To model temporal relationships between the various learning features, recurrent neural networks (RNNs) connect each time step with the previous ones through feedback connections. However, RNNs suffer the vanishing (and exploding) gradient problems, whereby the gradient of the error function becomes vanishingly small (or explosively large) in the long-term, preventing the backpropagation algorithm from updating the weights. This in turns prevents the RNNs from learning. Another option is the use of LSTM that we propose. Here, the computed weights can be ''forgotten'' thus keeping gradients unchanged. A simplified solution is known as the gated recurrent unit (GRU). GRUs simplify the forget gate operation through combining it with the input gate. This effectively allows for faster training of smaller datasets at a compromise of learning longer time sequences.

Often mentioned alongside RNNs are convolutional neural networks (CNNs). CNNs are useful in analyzing datasets that have grid patterns (e.g., classifying images or learning hierarchies in data) [17]. In comparison to RNNs (and by inheritence LSTMs), CNNs do not readily have temporal constructs. Instead, CNNs have a property of spatial shift invariance, which essentially means that small location translations do not impact the performance of CNNs. This is due to the pooling layer which effectively summarizes (i.e., through averaging or selection of extrema) the learning features. In the case of a set of pre-defined trajectories, CNNs may not offer any additional performance improvement, since these trajectories are splines in two-dimensional space and are—by definition—summaries.

1) COST FUNCTIONS

Are generally nonconvex functions of the true and predicted label vectors. Despite using functions that are seemingly convex (e.g., cross-entropy), the use of hidden layers with non-linear activation functions can lead to several ''equal'' global minima making the function nonconvex. As a result, an optimizer is needed to minimize the cost function.

2) BACKPROPAGATION

In order for the information to flow between layers, the backpropagation algorithm computes the gradient of the cost function with respect to the weights of a perceptron one layer at a time. Because of the nonconvexity of the loss function, a gradient-based optimizer is utilized to drive the prediction error to the lowest value possible.

3) TENSORS

Tensors are generalizations of two-dimensional matrixes to a larger dimension space. While the input data $\tilde{\mathbf{D}}$ is a two-dimensional time-indexed matrix, it has to be further reshaped to a tensor so that the third dimension is the time index.[1]

4) ACTIVATION FUNCTIONS

An activation function is a non-linear function[2] that defines the output of a perceptron from a given set of inputs and weights. Relevant ones in our paper are: 1) for LSTM, we use hyperbolic tangent as an activation for the memory cell state and sigmoid for recurrent input, forget, and output gates 2) for FCDNN, we use the rectified linear unit dense layers. The output layer of the model is a dense layer with a softmax activation function that calculates a probability for every possible class, which represents the unique beam identifier as motivated in (4).

5) INITIALIZATION

For both LSTM and FCDNN we initialize the bias to zeros and the weights according to the Glorot algorithm, such that the variance of the activations are the same across every layer. This helps prevent or minimize gradient saturation [19]. It an extensively used practice to optimize nonconvex functions though convergence may not be guaranteed [17].

6) OPTIMIZER

To find the cost optimal weights, we use the adaptive moments (Adam) optimizer [20]. Adam has benefits of being computationally efficient, uses the second moment to accelerate the descent, with little memory being required. Besides the learning rate which dictates the descent rate, other important parameters are 1) the batch size, which is the number of training samples in one optimizer step and 2) the number of epochs, which is the number of times the complete dataset is visited as part of the training phase.

7) HYPERPARAMETER TUNING

Hyperparameters are deep neural network settings that are used to control its behavior. To find the loss optimizing hyperparameter settings, we use grid search over a set of parameters related to the depth and width of the FCDNN. This is done over data samples that the deep neural network does not observe as part of its training, referred to as the ''validation'' data.

[1] The order of this dimension does not have to be the third. For example, with Keras [18], the time index is the first dimension, while the other two are simply a design matrix for each time index.

[2] An exception is the linear activation function, which is used when the output data type is continuous.

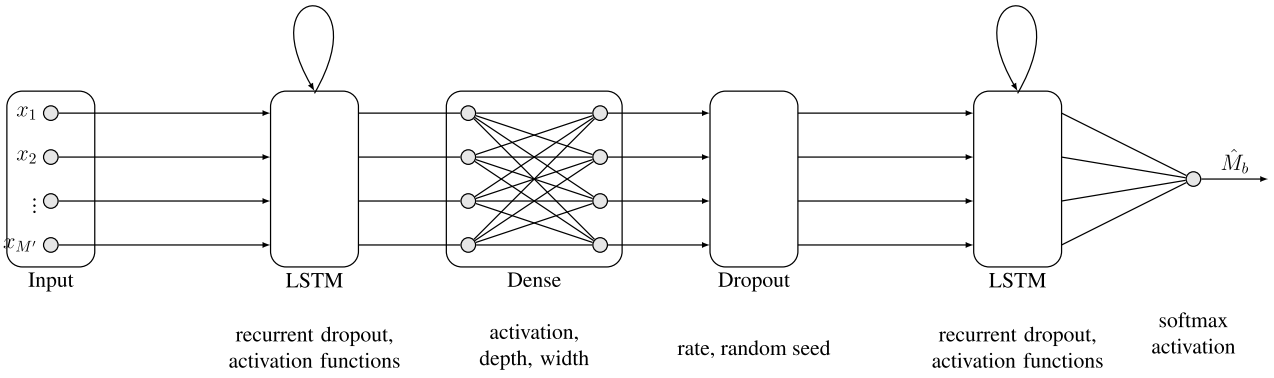

FIGURE 3. Architecture of the deep learner comprising fully connected deep neural networks, dropout layer, and LSTMs.

8) TRANSFER LEARNING

Transfer learning [6] is the process where a deep learner trained on one dataset is used 1) as a starting point for a model that performs a similar task or 2) to make inferences from a different and relevant dataset. Thus, transfer learning is more time-efficient than training from scratch. This becomes particularly useful in delay-sensitive RRM applications in wireless networks such as mobility.

E. LAYERS

Here, we focus on the ones that are relevant to our paper: 1) fully connected deep neural network (FCDNN), 2) long short-term memory (LSTM), and 3) dropout layers.

Fully connected deep neural networks are made of several layers of artificial neural networks to approximate a function [17]. The first layer is the "input layer" while the last layer is referred to as the "output layer." The hidden layers are all the layers between the input and the output layers. In a FCDNN, the number of nodes in each hidden layer is its width and the number of the hidden layers is its depth. FCDNNs do not factor in time relationships within the input data.

Long short-term memory are special types of RNNs. RNNs are a family of neural networks that can process sequential data and thus are suitable for time-indexed datasets. LSTMs comprise multiple elements: 1) memory cell state, which is generally common in RNNs 2) forget gate and 3) input gate [21]. The memory cell states simply record information. The input gate is used to update the memory cell state. The forget gate can learn to reset the state of the memory cell when the stored information is no longer needed (i.e., invalidation of the memory cell). If we denote vectors of the candidate memory cell (i.e., new information) as $\tilde{\mathbf{c}}_t$, the forget gate as $\mathbf{f}_t$, the input gate as $\mathbf{i}_t$, then we can represent the impact of the forget and input gate on the memory cell output as follows: $\mathbf{c}_t := \mathbf{c}_{t-1} \odot \mathbf{f}_t + \tilde{\mathbf{c}}_t \odot \mathbf{i}_t$.

Dropout layer randomly sets a fraction of the input units to zero at each training step, known as the *rate*. This effectively creates an ensemble of learners, which helps prevent overfitting (i.e., memorization of training data and failure of generalization). The dropout layers exist in both the LSTM and FCDNN layers.

Output layer selects the class of the highest likelihood. Since our deep learner is effectively a multi-class classifier, the softmax activation function in the output layer converts the values from the to a vector of real values the sum of which is equal to 1.

Architecture of the deep learner with the relevant hyperparameters per building block is shown in Fig. 3. A high-level reasoning for this architecture is explained as follows: we start with an LSTM to learn long occurring patterns, then we use a fully connected deep neural network to extract additional learning features from these occurring patterns. We drop some inputs at random to minimize overfitting and then we use another LSTM to learn any long term patterns from these additional features. These learnings are fed to the output layer, which uses the softmax activation function. The output of this function is a predicted class with the highest likelihood as motivated earlier.

F. TRAINING AND VALIDATION SPLIT

Given that $\mathbf{D}$ is a time-indexed dataset, a random split between a training and a validation set often performed in supervised learning is not feasible. This is because the implication of time and its impact on the values of each learning feature becomes overlooked. Instead, a split that has two types of constraints that have to be fulfilled: 1) no randomness and 2) split must abide by the time boundaries of a radio frame. To achieve these two constraints, the pivoting time index $t^\star$ at which the data is split is computed for a training data size r_{training} as follows:

$$t^\star := \lfloor \lfloor r_{\text{training}} m' \rfloor / N_{\text{slot}}^{\text{frame}} \rfloor N_{\text{slot}}^{\text{frame}}, \tag{6}$$

where the discrete interval $[0, t^\star - 1]$ is the time indexes belonging to the training data while the interval $[t^\star, n-1]$ is for the validation data. Here $N_{\text{slot}}^{\text{frame}}$ is the number of time slots in a radio frame.

G. RUN-TIME COMPLEXITY

Let us denote the LSTM input gate dimension as I and the memory cell state dimension as C. Further, let us denote the depth of the FCDNN as d and the width of it as w. Then, the run-time complexity of the architecture outlined

in Fig. 3 as affected by the change in these values is in $\mathcal{O}(IC + C^2) + \mathcal{O}(wd)$ [22], [23], where the first term is the run-time complexity of the LSTM and the second term is the run-time complexity of the FCDNN, both as a function of their respective dimensions.

V. BEAM HANDOFF ALGORITHMS

A. LEGACY

This follows the industry standards, where the network can configure UEs to measure and report the reference symbol for both serving and neighboring BS beams. For RRC-based mobility, the network can configures events for mobility decisions such that if the SS-RSRP of the neighboring BS is above the serving BS by a certain offset and time-to-trigger duration (i.e., A3 event) [3], the UE sends a measurement report to the network via the RRC protocol.

LLM on the other hand relies on beam-centric measurements where the network can configure UEs to periodically report serving and non-serving BS beams periodically [24].

The minimum report interval for RRC-based mobility is 120 ms [3] whereas the periodicity of the LLM beam reporting can be as low as 10 ms [12].

Because of the use of events and periodic reports, this algorithm does not benefit from any supervised learning and is essentially a reactive (as opposed to a proactive) approach.

B. PROPOSED

In our proposed algorithm, we use both LLM and RRC-based mobility procedures. The LLM procedure is adopted during the training for data collection where UEs report the SS-RSRP value per beam for several beams. There are other quantities construct the dataset and are either reported by the UE or computed at the BS as shown in Table 1.

Objective: The objectives tuple $\mathcal{O}$ are the radio measurements that the handoff aims at optimizing. This could be the beam reference symbol received power or any other target as reported per UE for all time indexes.

Trajectory: We define a trajectory of a given UE as a time-indexed tuple for the i-th UE $\mathcal{T}_i(t) := (t, i, \phi_i(t), b_i(t))$ collected over the period of movement of this UE in the association area of the network. Here, the elements of this tuple are the time index, C-RNTI, direction (ϕ) of the i-th UE, and beam identifier. Optionally a time-indexed tuple can contain the longitude and latitude coordinates $(t, i, \phi_i(t), x_i(t), y_i(t), b_i(t))$. The LSTM layers in the deep learner have an objective of learning these trajectories and the performance optimal beam as it changes over time over that trajectory. The union of the trajectory and objective tuples at a given time index t and for the i-th UE is essentially a row in the dataset $\mathbf{D}$. We can formally define $\mathbf{D}$ *before* the feature engineering steps as $\mathbf{D} := [\mathbf{X} \mid \mathbf{y}] = [(\mathcal{T}_i(t) \cup \mathcal{O}_i(t))_{i=1}^{u}]_t$.

Algorithm outline: Algorithm 1 is outlined as follows:

1) Collect UE LLM measurements and compute features as shown in Table 1 to construct $\mathbf{D}$. This contains the optimization objective measurement and helps construct trajectories of the UEs in motion.
2) Train a deep learner and use its predictions of objective-optimal beams for all UEs following similar trajectories as learned by the deep learner: If the predicted next optimal beam of the UE is different than the current serving beam, trigger a handover.
3) If a UE hands off to a wrong beam or BS (i.e., as a result of a poor deep learner prediction), which could lead to a handover failure as UE may experience a weak signal, invoke fallback: UE to initiate recovery procedures (e.g., RRC reestablishment or beam failure recovery). This effectively pauses the LLM measurements.
4) Perform transfer learning using new data (or *experience*) obtained from either good or bad predictions.
5) Resume LLM measurements.
6) This outline repeats as long as the algorithm is enabled.

Due to the predictions made by this algorithm, the UEs do not need to periodically report their radio measurements, which helps 1) conserve UE battery 2) reduce heat that is a major contributor to shot noise at mmWave frequencies and 3) reduce control signaling overhead. The algorithm is also robust against failure due to the fallback mechanism: If the prediction performance of the network is poor, the UE can either indicate beam failure or a RLF. Then, the UE and initiates an RRC re-establishment or beam failure recovery procedure to connect to the network. The proactive handoff for this UE is stopped until the re-establishment of the connection between UE and the network.

Training approaches: We propose two methods for model training: 1) Distributed and 2) Centralized.

1) DISTRIBUTED

In this case a deep learner is created per BS-UE pair. Thus, the UE identifier is necessary and transfer learning *cannot* be applied since the C-RNTI can be reassigned to a different UE [3]. This allows us to come up with two approaches of the distributed model:

Proposed with UE coordinates: In this scheme we allow the UEs to report its coordinates to the serving BS, which uses the coordinates as learning features to predict the optimal beam for these UEs.

Proposed without UE coordinates: UEs do not report their coordinates (x_i, y_i) to their serving BS.

In the distributed case, the number of deep learners required is $u|\mathcal{B}|$. This is because one deep learner is assigned per BS-UE pair.

2) CENTRALIZED

In this case a deep learner is created per BS and the BS does not require any knowledge about the UE identifier or positions. Instead, it ''crowdsources'' knowledge regardless of which UE reports it. This clearly reduces the number of deep learners required to $|\mathcal{B}|$.

Algorithm 1 Proactive Optimal Beam Handoff

Input: Optimization objective from $\mathcal{O}$, training-validation ratio from (6).
Output: Objective-optimal beam assignment to the i-th UE for the next radio frame for all $i \in \{1, 2, \ldots, u\}$.

```
Loop
  if moving UEs are in LLM then
    Collect reports and computations to construct D from the features shown in Table 1.
    Train a deep learner using first (6) samples and assess prediction performance:
    if validation performance is acceptable then
      Use deep learner predictions of optimal beams based on the trajectories as learned by the deep learner.
      Proactively handoff UEs to their predicted optimal beams if the predicted next best beam is different than the current serving beam.
      If UEs report un-degraded measurements (at least), set prediction to good.
    else
      Set prediction to bad.
    end
  end
  if prediction is bad then
    Fallback: UE initiates recovery procedure (e.g., RRC re-establishment or beam failure recovery).
    Instruct UEs to pause LLM measurements.
  end
  Retrain deep learner through transfer learning (if applicable) using recent experience from fallback.
  Instruct UEs to resume LLM measurements.
EndLoop
```

VI. SIMULATION

A. SETUP

The dataset used for the simulation is generated from a system level simulator based on the Manhattan grid as visualized in Fig. 4. This grid consists of two streets the width (and height) of each is 200 m (and 100 m). These distances are influenced by the empirical measurements from [25]. The radio propagation in this scenario provides street canyon like effects and sharp transitions between line of sight (LOS) and non-LOS. We use micro BSs each of which has an antenna height of 5 m. UEs are independently and uniformly distributed in the grid and move in one of four directions at random (uniform) and at an average speed of 100 km/h. We use full buffer traffic model (i.e., infinite data availability) to keep the system at maximum load. To add an element of complexity, the speed of the UEs changes every 100 m with probability of 0.2. We set r_{training} to 0.6 and run the simulation for a total of 50 seconds or 5,000 radio frames. Further details of the simulation setup can be found in Table 2.

The architecture of the deep learner is depicted in Fig. 3. The deep learner has depth of 8 and width of 16. It is trained with 192 epochs and a batch size of 64, each with a rectified linear unit activation function. The data training-validation split is 60%-40%. A dropout rate of 0.2 is used for the dropout layer. We also set the random generator seed to 0 for reproducibility.

We run three flavors of the proposed algorithm: 1) distributed with UE coordinates, 2) distributed without UE coordinates, and 3) centralized. The performance of these is benchmarked against a perfect predictor. The objective of interest is the SS-RSRP. That is, we proactively instruct UEs to handoff to received power-optimal beams as predicted by the next radio frame. For all three flavors (or approaches), we choose a lookback ℓ from $\mathcal{L} := \{0, 1, \ldots, 10\}$. Centralized approach ignores two features from: 1) C-RNTIs (which are used in the distributed approach) and 2) the UE coordinates.

B. RESULTS

We start this subsection with a few important definitions related to the performance of the proactive beam handoff algorithm. These are the accuracy and the relative accuracy.

1) AVERAGE ACCURACY

The average accuracy for a multi-class classification can be challenged as a suitable measure since the beam idenitifiers may not be balanced in their count. However, what the UE is interested in a proactive handoff is whether the predicted and true beam identifier are identical. Thus, the accuracy ($\bar{A}$) is defined as an average of identical values of the classes:

$$\bar{A}(\mathbf{y}, \hat{\mathbf{y}}) := \frac{1}{V} \sum_{i=0}^{V-1} \mathbb{1}\left[[\mathbf{y}]_i = [\hat{\mathbf{y}}]_i\right], \tag{7}$$

where V is the dimension of the beam identifier vector, collected over the validation time period (6) for all UEs. The average accuracy is a significant measure for radio access performance since a mispredicted beam results in wrong handoff decision and thus can degrade quality of experience and lead to mobility failure or increase the duration of service interruption [26]. Also, the RLF rate due to beam handoff failure is simply $1 - \bar{A}$.

2) ZERO-ONE SCORE

To assess the added value of the UE position to the prediction performance, we define a vector of loss-win scores known as the zero-one score ($\boldsymbol{\alpha} \in \{0, 1\}^V$) as:

$$\boldsymbol{\alpha}(\hat{\mathbf{s}}, \mathbf{s}_{\hat{\text{base}}}; \varepsilon) := \mathbb{1}\left[(\hat{\mathbf{s}} - \mathbf{s}_{\hat{\text{base}}}) > \varepsilon \mathbf{1}\right], \tag{8}$$

where the V-dimensional column vectors are for a certain performance measure (e.g., average accuracy), ε is the cutoff threshold, and $\mathbb{1}[\cdot]$ is overridden as the indicator *vector*.

TABLE 1. Dataset features.

Parameter	Data type	Owner
Radio frame number	Integer	Computed by the BS
C-RNTI	Integer	Assigned by BS to the UE
Current beam index	Integer	Reported by the UE*
Previous beam index	Integer	Reported by the UE*
Beam SS-RSRP	Float	Reported by the UE
UE direction	Integer	Computed by the BS
UE speed	Float	Computed by the BS
UE coordinates	Float	Reported by the UE
Radio link failure indicator	Boolean	Reported by the UE

* These are assigned by the serving BS.

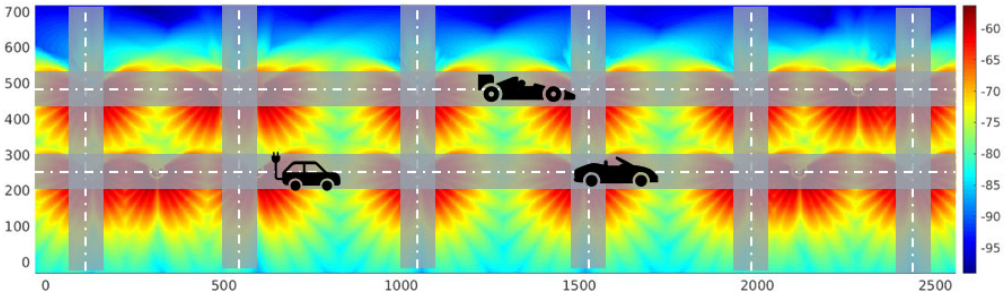

FIGURE 4. Manhattan grid. Base stations are uniformly distributed at each road intersection. The received signal power per beam [dBm] is displayed for all beam candidates.

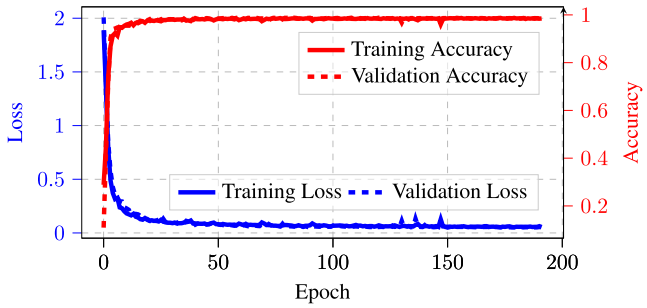

FIGURE 5. Accuracy and loss for both the training and validation data for a given BS-UE pair and $\ell = 7$ lookback.

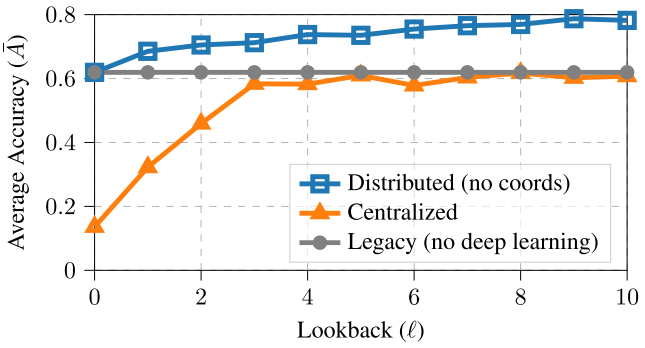

FIGURE 6. The average accuracy of the proactive beam handoff approach vs. lookback.

For the results, we begin with the loss (5) as shown in Fig. 5. As the number of epochs increase we observe a decrease in the loss and increase in the accuracy which shows that the deep learner is indeed "learning" from **D**. This demonstrates that proactive beam handoff using deep learning is possible.

In Fig. 6, we show the average accuracy of all legacy, distributed, and centralized approaches over several lookback values. The performance of the distributed approach to proactive handoff is superior to the centralized one. With $\ell = 0$, the distributed case with no coordinates is equivalent to the legacy approach outlined in Section V-A. We discuss why the performance of the distributed approach is superior to the centralized one in the next subsection.

Fig. 7 shows the zero-one score with parameter $\varepsilon = 0.03$. Here, the baseline is the performance of the distributed approach without UE coordinates. We observe that the distributed approach using the UE coordinates does not always outperform baseline. Particularly, we observe that the deep learner does not perform well with short lookback sequences, which is expected as LSTM generally penalizes shorter sequences. However, when the number of lookbacks is between 2 and 7, then the distributed approach with UE coordinates outperforms the baseline. We also notice that for larger ℓ, knowledge of UE coordinates does not bring any additional insight to the deep learner. Further details are discussed in the next subsection.

To understand the extent of the performance of the distributed approach using the UE coordinates, we plot the aggregated accuracy values for all BS-UE pairs grouped by the lookback values in Fig. 9. The takeaway is that the performance of the model varies across different users in the network and that not all users would similarly benefit from a proactive handoff.

TABLE 2. Simulation settings.

Technology and duplex mode	5G TDD
Channel model	3GPP TR 38.901 Urban Micro [27]
Beams per base station (BS)	8
BS azimuth beam angles	$\{44°, 56°, 69°, 82°, 96°, 109°, 122°, 134°\}$
Traffic model	Full Buffer
Timeslots per radio frame	10
BS transmit power	44 dBm
Number of BSs ($\lvert\mathcal{B}\rvert$)	32
Number of UEs (u)	400

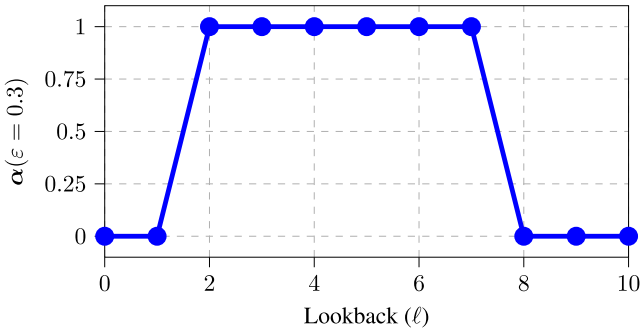

FIGURE 7. Vector plot of α of the difference between impact of UE positions compared to without UE positions for the distributed model.

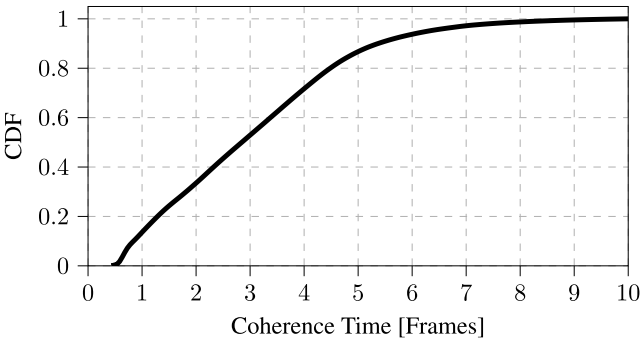

FIGURE 8. Channel coherence cumulative distribution function (CDF).

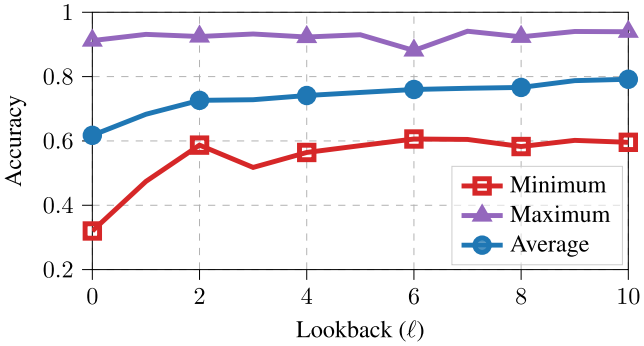

FIGURE 9. The maximum, average, and minimum prediction accuracy of the distributed proactive beam handoff algorithm with various lookbacks.

C. DISCUSSION

We start with discussing the performance disparity in both the centralized and distributed approaches with no UE coordinates in Fig. 6. This disparity is attributed to a wide range of trajectory-dependent measurement values that several UEs simultaneously take. In other words, a given point on the trajectory can have several different measurements depending on which serving beam they were on. This applies to all points on the trajectory alike. The absence of the C-RNTI drives the centralized deep learner (one learner per BS) to learn a weighted average of measurements which prevents accurate prediction of the objective-optimal beam identifier. However, with the distributed approach, the C-RNTI allows these trajectory-dependent measurements to be uniquely attributed the UE reporting them, causing the learner to learn unique measurements along the trajectory—a task that an LSTM is designed to handle. A natural question after this discussion is: How do these two algorithms compare with the legacy algorithm where no proactive beam handoff exists? To answer this, we remark that the legacy algorithm is a per BS-UE algorithm. This is because each UE either reports an RRC event or a periodic beam measurement independently from other UEs, and the BS decides whether or not a beam handoff should take place. Thus, the comparison with the distributed method of our proposed algorithm is more suitable (which is also BS-UE pair). It is clear that the legacy algorithm does not depend on the number of lookbacks since it depends on the most recent measurement only. We observe that the legacy approach underperforms the distributed proposed approach. One reason why this happens is because the knowledge of the trajectory and its underlying radio conditions during the handoff enables the proactive approach to assign a more robust beam (e.g., compared to a beam with stronger SS-RSRP that is short-lived in the trajectory that the legacy algorithm assigns).

Next, we study the impact of the knowledge of UE coordinates in the proactive beam handoff. We do this through studying the impact of lookback on the zero-one score as shown in Fig. 7. Smaller sequences could drive LSTMs towards over-triggering their forget gates, which deletes the information in the self-recurrent memory BSs. However, as the number of sequences grown larger, we observe two important behaviors which can be mapped to two different regimes:

1) Short lookback regime: any number of lookbacks below 7 (but greater than 1) radio frames.
2) Long lookback regime: any number of lookbacks greater than 7 radio frames.

An important question can be asked here: What is the significance of such a breakdown? To answer this question, we observe that the win-loss score for the distributed approach using the UE coordinates versus the approach not using them is equal to 0 after $\ell = 7$. This is because of the capability of LSTM to learn long sequences. As a result, the knowledge of UE coordinates is no longer relevant since the deep learner can implicitly learn them with high accuracy through the beams and the time-based trajectory. At $\ell = 7$, we have a total of $n = 2(7 + 1 + 1) = 18$ learning features for $\mathbf{D}$. We cannot extrapolate ℓ indefinitely as ℓ cannot exceed the beam coherence time (3). We show the coherence time for the UEs in Fig. 8, where only 1.2% of the samples represent coherence times greater than 7 radio frames. This is expected for high-mobility UEs.

Finally, we notice that not all UEs benefit from the proactive beam handoff equally. Fig. 9 shows that while some UEs have an accuracy consistently above 0.9 for any lookback value, there are some UEs that do not enjoy that. This depends on the proximity of these UEs to their serving BS and the uniqueness of the trajectories they follow. Based on the Manhattan grid in Fig. 4, we compute that the inter-site distance can be as short as 100 m and as long as 460 m. If trajectories have small number of overlaps, then LSTMs can learn the trajectory and the optimal beams on its path. However, if UEs are on trajectories that overlap with other trajectories, such as turns, then undesired handoff effects such as ping-pong could cause conflicting data to be added to the dataset. This would cause the LSTMs to fail in learning the optimal beams. If beams are equally likely to serve UEs, then the lower bound of accuracy is 0.125, which is the reciprocal of the number of beams 8, assuming equally likely occurrence of beams.

VII. CONCLUSION

In this paper, we demonstrated that proactive beam handoff is possible using deep learners with deep learners created per BS-UE pair residing at the edge. We showed three different approaches. The introduction of the UE coordinates help enhance the performance of the deep learner towards the proactive beam handoff; however as the longer sequences are presented to the deep learner, the need for the UE coordinates becomes diminished. The use of intelligence in radio resource management algorithms help reduce the dependency on the UEs therefore preserving their battery life and reducing their temperature, which helps combat shot noise—a major issue in mmWave communications.

REFERENCES

[1] A. Alkhateeb, I. Beltagy, and S. Alex, "Machine learning for reliable mmWave systems: Blockage prediction and proactive handoff," in *Proc. IEEE GlobalSIP*, 2018, pp. 1055–1059.

[2] *NR; Physical Layer Procedures for Control*, document TS 38.213, 3rd Generation Partnership Project, Jul. 2020.

[3] *NR; Radio Resource Control (RRC); Protocol Specification*, document TS 38.331, 3rd Generation Partnership Project, Jul. 2020.

[4] *Revised SID: Study on Enhancement for Data Collection for NR and ENDC*, document SID Revised RP-201620, 3rd Generation Partnership Project, Sep. 2020.

[5] *New SI: Study on Artificial Intelligence (AI)/Machine Learning (ML) for NR Air Interface*, document SID New RP-213599, 3rd Generation Partnership Project, Dec. 2021.

[6] 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.

[7] F. B. Mismar and J. Hoydis, "Unsupervised learning in next-generation networks: Real-time performance self-diagnosis," *IEEE Commun. Lett.*, vol. 25, no. 10, pp. 3330–3334, Oct. 2021.

[8] F. B. Mismar, A. Alammouri, A. Alkhateeb, J. G. Andrews, and B. L. Evans, "Deep learning predictive band switching in wireless networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 96–109, Jan. 2021.

[9] F. B. Mismar and B. L. Evans, "Deep learning in downlink coordinated multipoint in new radio heterogeneous networks," *IEEE Wireless Commun. Lett.*, vol. 8, no. 4, pp. 1040–1043, Aug. 2019.

[10] J. Zhao, J. Liu, Y. Nie, and S. Ni, "Location-assisted beam alignment for train-to-train communication in urban rail transit system," *IEEE Access*, vol. 7, pp. 80133–80145, 2019.

[11] A. Namdari, M. A. Samani, and T. S. Durrani, "Lithium-ion battery prognostics through reinforcement learning based on entropy measures," *Algorithms*, vol. 15, no. 11, p. 393, Oct. 2022. [Online]. Available: https://www.mdpi.com/1999-4893/15/11/393

[12] *NR; Medium Access Control (MAC) Protocol Specification*, document TS 38.321, 3rd Generation Partnership Project, Jun. 2022.

[13] *NR; NR and NG-RAN Overall Description*, document TS 38.300, 3rd Generation Partnership Project, Jul. 2020.

[14] D. Chizhik, "Slowing the time-fluctuating MIMO channel by beam forming," *IEEE Trans. Wireless Commun.*, vol. 3, no. 5, pp. 1554–1565, Sep. 2004.

[15] V. Va, J. Choi, and R. W. Heath Jr., "The impact of beamwidth on temporal channel variation in vehicular channels and its implications," *IEEE Trans. Veh. Technol.*, vol. 66, no. 6, pp. 5014–5029, Jun. 2017.

[16] *Moderator Summary for Multi-Beam Enhancement: EVM, 3GPP TSG RAN WG1 #102-e e-Meeting*, document R1-2007151, 3rd Generation Partnership Project, Aug. 2020.

[17] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: http://www.deeplearningbook.org

[18] F. Chollet. (2015). *Keras*. [Online]. Available: https://github.com/fchollet/keras

[19] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. Int. Conf. Artif. Intell. Statist.*, vol. 9, May 2010, pp. 249–256.

[20] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. for Learn. Represent.*, May 2015, pp. 1–15.

[21] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[22] M. Rotman and L. Wolf, "Shuffling recurrent neural networks," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, no. 11, pp. 9428–9435.

[23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Nov. 2011.

[24] *LS on Agreements Pertaining to L1/L2-Centric Inter-Cell Mobility*, document R1-2102209, 3rd Generation Partnership Project, Feb. 2021.

[25] S. Sun, G. R. MacCartney, and T. S. Rappaport, "Millimeter-wave distance-dependent large-scale propagation measurements and path loss models for outdoor and indoor 5G systems," in *Proc. 10th Eur. Conf. Antennas Propag. (EuCAP)*, Apr. 2016, pp. 1–5.

[26] A. O. Kaya and H. Viswanathan, "Deep learning-based predictive beam management for 5G mmWave systems," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2021, pp. 1–7.

[27] *5G; Study on Channel Model for Frequencies From 0.5 to 100 GHz*, document TR 38.901, 3rd Generation Partnership Project, Nov. 2020.

FARIS B. MISMAR (Senior Member, IEEE) was born in IL, USA. He received the B.S. degree (Hons.) in computer engineering from The University of Jordan, in 2004, the M.S. degree (Hons.) in electrical engineering and the M.B.A. degree (Hons.) from The University of Texas at Dallas, in 2011 and 2014, respectively, and the Ph.D. degree in electrical and computer engineering from The University of Texas at Austin, in 2019.

From 2004 to 2008, he was a Telecommunications Engineer with Motorola. From 2009 to 2018, he was a Senior Solutions Manager with Radio Access Networks, Ericsson. From 2019 to 2020, he was the Director of radio performance assurance and engineering with Samsung Electronics America. He joined Nokia in 2020, as the Head of the Data Science and Analytics Practice, Mobile Networks. Since 2022, he has been a Senior Principal Consultant with Nokia Bell Labs Consulting, Dallas, TX, USA. Since 2021, he has been an Adjunct Assistant Professor with the Electrical Engineering and Computer Science Department, The University of Texas at Dallas. He is the author of more than ten conferences and journals. He served as a Consultant and an Advisor for many leading wireless communications service providers and governments in different continents. He holds two patents with several pending. His research interests include radio resource management in cellular communications using machine learning.

Dr. Mismar was a recipient of the University of Jordan Award for Excellence, in 2004, the Marcus Wallenberg Foundation Scholarship for Scientific Research and Education Scholarship Award, in 2016 and 2017, and the Nokia Bell Labs Distinguished Member of Technical Staff Recognition Award, in 2023.

ALPEREN GÜNDOĞAN (Member, IEEE) was born in Malatya, Turkey. He received the B.Sc. degree in electronics and communication engineering from Istanbul Technical University, in 2017 and the M.Sc. degree in communication engineering from the Technical University of Munich, in 2020.

He was a Radio Research Specialist with Nokia Bell Labs, from 2020 to 2022. He joined Apple, as a Wireless Researcher Engineer, in 2022. He has contributions on many 3GPP technical documents in RAN1/RAN2/RAN3 with system level simulation results and concept work. He is the author of two conference papers and inventor of more than 20 filed patent applications.

ALIYE ÖZGE KAYA (Senior Member, IEEE) received the B.Sc. and Dipl-Ing. degrees in electrical and information technology engineering from the Technical University of Munich, Germany, in 2003, the M.Sc. and Ph.D. degrees in electrical and computer engineering from the Wireless Information Networks Laboratory (WINLAB), Rutgers University, NJ, New Brunswick, USA, in 2010, and the M.B.A. degree from the Columbia Business School. She has been a Research Scientist with Nokia Bell Labs, since 2010. Her current research interests include applying artificial intelligence and machine learning techniques to wireless communications systems. She was a recipient of the Best Paper Award from the Next Generation Networking and Internet Symposium of IEEE GLOBECOM 2009 for work related to P2P-ISP cooperation.

OLEG CHISTYAKOV (Member, IEEE) was born in Saint Petersburg, Russia. He received the M.S. degree in applied mathematics from Saint Petersburg Electrotechnical University "LETI," in 1995.

From 1995 to 1998, he was a DSP Engineer with Digiton. From 1998 to 2009, he was a Senior DSP Software Design Engineer with Lake Communications. From 2009 to 2011, he was a Senior Software Design Engineer with DA-Design Electronics. From 2011 to 2014, he was a Senior Software Engineer with Airspan Networks. Since 2014, he has been a Senior Research Engineer with Nokia Bell Labs.