

Deep Learning-Based mmWave Beam Alignment with Only Pilot Channel Measurements

Taekyun Lee, Hyeji Kim, and Jeffrey G. Andrews

6G@UT, Dept. of ECE, The University of Texas at Austin, USA

Email: taekyun@utexas.edu, hyeji.kim@austin.utexas.edu, jandrews@ece.utexas.edu

Abstract—For millimeter wave (mmWave) communication, fast and accurate beam alignment is essential but challenging. Site-specific beam adaptation using deep learning is a very promising paradigm for beam alignment, but such methods typically require a lot of clean channel measurements for training, which can be difficult or even impossible to achieve in practice. This paper introduces a novel method to learn beam alignment policies using only uplink (UL) pilot measurements. The proposed method integrates a generative adversarial network (GAN)-based channel estimation (CE) model with an unsupervised deep learning model beam alignment engine (BAE). We introduce an efficient form of dataset amplification for improved training that leverages the randomness of the deep generative model (DGM) and an early stopping mechanism. Our experiments show that the GAN-BAE method achieves a better signal-to-noise ratio (SNR) by nearly 3 dB compared to compressed sensing (CS) methods such as orthogonal matching pursuit (OMP) and EM-GM-AMP (an Approximate Message Passing algorithm), especially when there are limited pilot measurements from each mobile user.

I. INTRODUCTION

Beam alignment is a key bottleneck in millimeter wave (mmWave) systems, usually requiring large bandwidth, time slot, and energy overheads [1]. As cellular systems continue to migrate towards high carrier frequencies, fast and efficient beam alignment will become increasingly important. This underscores the need for more resilient and adaptive mmWave beam codebook designs and beam alignment algorithms [2].

Site-specific beam adaptation is a promising approach to beam alignment. The basic idea is that each cell in the network learns a site-specific codebook-based (CB) that is efficient and effective for the propagation conditions and channels commonly experienced in that cell – thus it depends on the UE distribution also. By concentrating the initial access (probing) beam searches such that they are aligned with the learned channel distribution in the cell, many fold gains can be observed in terms of beam alignment latency [3]. Conventionally, most beam management approaches including the 5G NR initial access procedure assume a CB-based method [3]–[5]. Our recent work [6] was able to find a near-optimal beam pair within the continuous angular domain, resulting in better gain compared to the CB method. The results of this paper apply to any deep learning-based BAE approach.

This work was partly supported by NSF Award CNS-2148141, as well as ARO Award W911NF2310062, ONR Award N00014-21-1-2379, NSF Award CNS-2008824, and Keysight Technologies through the 6G@UT center within the Wireless Networking and Communications Group (WNCG) at the University of Texas at Austin.

Training a deep learning-based beam alignment engine (BAE) requires a fairly large channel dataset, on the order of 100K channel measurements for a typical macrocell sector [3], [6], [7]. Instead of generating a large synthetic dataset using statistical models, ray tracing, and/or simulations, we prefer an approach that could result in automated training of the BAE in a given cell just using existing over-the-air channel measurements like UL pilot measurements, such as Sounding Reference Symbols (SRS) or Demodulation Reference Symbols (DM-RS) in 5G.

Advances in channel estimation (CE) techniques allow for the creation of highly accurate underlying channels using pilot measurements. Such examples include linear regression methods such as least square (LS) and minimum mean square error (MMSE), compressed sensing (CS) methods such as orthogonal matching pursuit (OMP) [8] and EM-GM-AMP (an Approximate Message Passing algorithm) [9]. Additionally, there are discriminative deep neural network (DNN)-based CE [10], [11], and deep generative model (DGM)-based CE such as generative adversarial network (GAN) and diffusion models [12], [13].

However, downlink (DL) CE is inapplicable as an input to the BAE for two key reasons: (a) The *limited* volume of pilot measurements, which will be on the order of the number of users, is insufficient as BAE requires more extensive channel data for training. (b) We anticipate that the BAE will be trained at the BS, which has more computing power than a UE. This implies that the BS needs to access the estimated channels. In the traditional DL CE framework, however, the BS sends DL pilots, based on which the UE estimates the channel. Requiring UEs to transmit an estimated channel to the BS would introduce a substantial overhead.

We propose a novel framework that addresses these two challenges in automating the training of the BAE using pilot channel measurements. To be concrete, this paper builds on recent advancements in two parallel lines of research. First, the GAN-based framework for high dimensional CE [12] using UL pilots. Second, the end-to-end grid-free deep learning for site-specific mmWave beam alignment [6], which we use as our BAE. The contributions of this paper are as follows:

Beam alignment training system requiring only over-the-air UL measurements. We can train the GAN-BAE using only uplink (UL) pilot measurements, without any prerequisite dataset. Most of the GAN and the BAE training takes place at the base station (BS), as desired, given the substantial power




Fig. 1: The structure and process of the proposed method.

requirements of the DGM. The framework is a novel approach that can be carried out **completely autonomously over-the-air** by leveraging the downlink-uplink channel reciprocity which is completely new in the area of deep learning-aided beam alignment approaches, to the best of our knowledge.

Amplifying channel dataset using a DGM. We harness the *generative* capability of DGM-based channel estimation, which enhances dataset robustness, by generating an amplified channel dataset from minimal pilot measurements. Our method can outperform beamforming (BF) that are using a CS CE with the same number of measurements. In addition, to compensate for the increased processing time from the above steps, we propose a method to reduce the epoch during the GAN inference.

Beam alignment SNR roughly doubled. Applying the two methods mentioned above to train the BAE, the GAN-based results exhibit better BF signal-to-noise ratio (SNR) by 2.8 dB than CS-based results, achieving almost twice the gain. Although GAN-based CE show higher normalized mean square error (NMSE) values than CS methods in CE, it might provide a better estimate for the BF especially when the end goal of the channel estimate is simply to find a single optimal beamforming direction.

II. SYSTEM MODEL

We consider a single user downlink (DL) narrowband mmWave multiple-input multiple-output (MIMO) communication scenario, where the BS has an array of N_B antennas and the user equipment (UE) has an array of N_U antennas, each performing beam alignment. The BS and the UE sweep N_{probe} independent probing beams each, where each probing beam consists of N_{pilot} symbols.

We assume MIMO systems operating within a single-cell time division duplex (TDD) so that the channel remains constant within each time slot. Due to the channel matrix reciprocity, it can be inferred that the UL pilot signal grants the BS an estimation of the forward channel in the DL as well [14].

Fig. 1 provides an overview of the system model. First, we train the GAN generator using a significant amount of accumulated UL pilot measurements at the BS. Next, we use

the GAN generator to predict the channel from real-time UL pilot measurements. We then use these estimated channels to train the DL probing beam structure of BAE. Finally, we select a DL beam pair. Each step is detailed in Section III-A.

We will describe the formulas of the pilot measurements, probing measurements, and final received signals. Denote the hybrid precoder at the transmitter by $\mathbf{F}_{\text{UL}} \in \mathbb{C}^{N_U \times N_{\text{probe}}}$, and the hybrid combiner at the receiver by $\mathbf{W}_{\text{UL}} \in \mathbb{C}^{N_B \times N_{\text{probe}}}$. With the DL MIMO channel denoted by $\mathbf{H} \in \mathbb{C}^{N_U \times N_B}$, the transmitted pilot symbols $\mathbf{s}_{\text{UL}} \in \mathbb{C}^{N_{\text{probe}} \times N_{\text{pilot}}}$ are received as the composite received signal $\mathbf{Y}_{\text{UL}} \in \mathbb{C}^{N_{\text{probe}} \times N_{\text{pilot}}}$. The UL received signal can be written as

$$\mathbf{Y}_{\text{UL}} = \sqrt{P_U} \mathbf{W}_{\text{UL}}^H \mathbf{H} \mathbf{F}_{\text{UL}} \mathbf{s}_{\text{UL}} + \mathbf{W}_{\text{UL}}^H \mathbf{n}, \quad (1)$$

where P_U is the transmit power of UE, and $\mathbf{n} \sim \mathcal{CN}(0, \sigma^2 \mathbf{I})$ is the UL measurement noise.

After receiving the pilot measurements, BS sweeps N_{probe} probing beams while the UE measures the received signal power using N_{probe} probing beams. Let the BS probing beam pairs $\mathbf{F} \in \mathbb{C}^{N_B \times N_{\text{probe}}}$ and the UE probing beam pairs $\mathbf{W} \in \mathbb{C}^{N_U \times N_{\text{probe}}}$ to transmit the probing symbols $\mathbf{s} \in \mathbb{C}^{N_{\text{probe}} \times N_{\text{pilot}}}$. The DL received signal of all combinations of probing beams can be written as

$$\mathbf{Y} = \sqrt{P_B} \mathbf{W}^H \mathbf{H} \mathbf{F} \mathbf{s} + \mathbf{W}^H \mathbf{n}, \quad (2)$$

where $\mathbf{s}$ is the vector of transmitted signals and $\mathbf{n} \sim \mathcal{CN}(0, \sigma^2 \mathbf{I})$ is the DL noise.

The DL beam synthesizer is trained as detailed in Section III-B. Finally, the DL signal is transmitted from the BS to the UE. The BS selects a synthesized beam $\mathbf{f} \in \mathbb{C}^{N_B \times 1}$, and the UE selects a synthesized beam $\mathbf{w} \in \mathbb{C}^{N_U \times 1}$ to transmit a symbol $s \in \mathbb{C}$. In order to reduce the cost and power consumption of a fully digital system, analog BF is assumed in this work. The DL received signal can be written as

$$y = \sqrt{P_B} \mathbf{w}^H \mathbf{H} \mathbf{f} s + \mathbf{w}^H \mathbf{n}, \quad (3)$$

where P_B is the BS power and $\mathbf{n} \sim \mathcal{CN}(0, \sigma^2 \mathbf{I})$ is the DL noise. The SNR achieved by the chosen DL beam pair is

$$\text{SNR} = \frac{P_B |\mathbf{w}^H \mathbf{H} \mathbf{f}|^2}{|\mathbf{w}^H \mathbf{n}|^2}. \quad (4)$$




Fig. 2: The training phase of the proposed GAN-based CE model.




Fig. 3: The proposed GAN-BAE model, described in Section III-C.

III. PROPOSED DNN MODEL

A. GAN-based CE

We introduce the GAN-based CE model [12]. The beamspace representation of mmWave MIMO channel has a high spatial correlation. Therefore, we can make discrete Fourier transform (DFT) transformation to exploit the clustering in the angular domain directly as

$$\mathbf{H} = \mathbf{A}_R \mathbf{H}_v \mathbf{A}_T. \quad (5)$$

When assuming uniformly spaced linear arrays at both the transmitter and receiver, both $\mathbf{A}_R$ and $\mathbf{A}_T$ are unitary matrices. As a result, $\mathbf{H}$ and $\mathbf{H}_v$ share the same dimension. UL pilot measurement can be expressed as

$$\mathbf{Y}_{UL} = \mathbf{A} \mathbf{H}_v^H + \tilde{\mathbf{n}}, \quad (6)$$

where $\mathbf{A}$ is defined in [12]. Therefore, recovering $\mathbf{H}_v$ becomes an inverse problem, and a DGM $\mathbf{G}$ can be used to solve it approximately [15]. Specifically, a Wasserstein GAN with gradient penalty (WGAN-GP) is deployed in this work.

During the training process, we cannot access pure $\mathbf{H}$ channels but only noisy pilot measurements. Consequently, directly incorporating $\mathbf{H}_v$ is not feasible. To address this issue, we utilize the LS estimate from pilot measurements to generate the inputs for the GAN.

We assume to have $K = \lceil N_U / N_{\text{pilot}} \rceil$ precoder and combiner pairs, which make K different pilot measurements for each

UE. This ensures sufficient rank of stacked $\mathbf{A}$. Then, K pre-train pilot measurement equations can be constructed, and the beamspace LS channel estimate $\mathbf{H}_{v,LS}$ can be obtained as shown in Fig. 2.

During inference, we use $\mathbf{G}$ to output samples $\mathbf{H}_v$, by optimizing a latent vector using stochastic gradient descent with inference pilot measurements $\mathbf{Y}_{UL,test}$ as

$$\mathbf{z}^* = \arg \min_{\mathbf{z} \in \mathbb{R}^d} \left\| \mathbf{Y}_{UL,test} - \mathbf{A} \mathbf{G}(\mathbf{z}) \right\|_2^2 + \lambda_{\text{reg}} \|\mathbf{z}\|_2^2. \quad (7)$$

Therefore, by training $\mathbf{G}$ and the subsequent optimization in (7), the beamspace channel estimate is given by $\mathbf{H}_{v,est} = \mathbf{G}(\mathbf{z}^*)$, as shown in Fig. 3.

B. The grid-free beam alignment engine (BAE)

The BAE introduced in [6] aims to learn N_{probe} transmit (Tx) probing beams and N_{probe} receive (Rx) probing beams that are tailored to the overall channel distribution in the cell (and thus site-specific). Site-specific probing beams will effectively obtain the pilot measurements sent from UE to BS. After the probing beam sweeping, the BS uses collected measurements as inputs to the transmitter beam synthesizer function f_{BS} and receiver beam synthesizer function f_{UE} .

We use DL probing beam measurements in (2) as an input to the beam synthesizer function

$$\mathbf{Y}_{all} = \left[|\text{diag}(\mathbf{Y})|_1|^2 \quad \cdots \quad |\text{diag}(\mathbf{Y})|_{N_{\text{probe}}}|^2 \right]^T. \quad (8)$$

The beam synthesizer functions should generate beams that maximize the BF SNR for each channel realization. The probing beams should serve two parallel purposes simultaneously. First, they must provide well-chosen pilot signals as input to the beam synthesizer function. Second, they should be able to discover new UEs during the initial access (IA) process.

The BAE optimizes $\mathbf{F}$, $\mathbf{W}$, $\mathbf{f}$, and $\mathbf{w}$ using a complex neural network (NN) module. For $\mathbf{F}$ and $\mathbf{W}$, the complex NN module computes the composite matrix of received signals, sweeping probing beam pairs in a differentiable manner. For $\mathbf{f}$ and $\mathbf{w}$, the multilayer perceptron (MLP) uses two hidden layers with rectified linear unit (ReLU) activation and a final linear layer normalizing element-wise to enforce the unit-modulus constraint for analog BF.

Beam synthesizers f_{BS} at the BS and f_{UE} at the UE sweep probing beam pairs periodically. The UE measures and reports received signal power vector $\mathbf{Y}_{all}$ to the BS. Using these measurements, the BS and UE adjust the Tx and Rx beams. Because these beams are site-specific, the BS transmits components $\mathbf{W}$ and f_{UE} to the UEs, possibly through a low-frequency side link as supported in 5G systems. The BAE model is illustrated in Fig. 3.

However, training the BAE for each site demands about 100K of clean channel matrices, which is nearly impossible. This challenge leads us to the method discussed in Section III-C, where we leverage GAN to estimate channels from pilot measurements.

TABLE I: Simulation Parameters

Scenario Name	Outdoor LOS (O1)
UE training samples	467265
UE inference attempts	46
Number of pilot symbols (N_{pilot})	8
BS Antenna	64×1 ULA
UE Antenna	16×1 ULA
GAN train epochs	30000
GAN inference epochs	50, 500
BAE probing beams (N_{probe})	16
BAE training epochs	1000
Carrier frequency	28 GHz
Bandwidth (B)	100 MHz
BS power (P_B)	30 dBm
UE power (P_U)	-10, 0, 10, 20 dBm
Noise power (σ^2)	-81 dBm




Fig. 4: Illustration of the DeepMIMO ray-tracing scenarios.

C. Proposed unified model (GAN-BAE)

In this section, we introduce a novel framework that integrates the GAN-based channel estimation and the grid-free beam alignment technique.

Setup. For concreteness, we consider the DeepMIMO dataset [7] built from a ray-tracing-based channel simulator. In this work, a 28 GHz outdoor scenario is considered. Table I shows the exact parameters and settings. The O1 scenario is an outdoor urban simulation environment with line-of-sight (LOS) paths of UEs.

Our framework: GAN-BAE. Our framework integrates the GAN-based CE from Section III-A and the BAE from Section III-B. The unified model, GAN-BAE, operates as follows: We collect UL pilot data at the BS over an extended period, allowing the GAN to capture the channel matrix distribution. This process is for the dataset generation before the actual transmission, which only needs to be performed once in the same environment if there are no significant channel distribution shifts due to environmental changes. Once the GAN generator is trained, we periodically collect a limited set of real-time UL pilot measurements and amplify the CE dataset from the pilot measurements using the GAN generator to train the BAE. Consequently, the BAE learns the DL probing beam and beam synthesizer, which can be conceptually viewed as infinitely large codebooks.

In our experiment, the GAN is trained using a dataset of 467,219 UE instances, which are assumed to have been collected in advance. The dataset is split into training and

validation sets at a 3:1 ratio. This division helps determine when to halt the GAN training.

At each time instance, we collect extremely few pilot measurements. Specifically, we assume that 46 samples are sparsely collected from the user grid, as shown in Fig. 4. This represents 0.01% of the UEs. There are two reasons for using such a limited number of pilot measurements. First, we aim to align with the current channel distribution, so outdated data is inappropriate. Second, training the BAE necessitates feedback from the UE to the BS, which is time-consuming with real-time UL measurements. While previous study [6] has assumed offline training with synthetic channel datasets, the GAN-BAE seeks to utilize real-time pilot measurements for online BAE training. Hence, we must achieve effective BF with minimal UL pilot measurements.

Handling limited pilot measurements presents a challenge in developing a high-quality dataset for training the BF system. While data augmentation emerges as a potential solution, traditional methods tend to enhance robustness in the image domain, often neglecting the unique characteristics of our desired channel-like output. Instead of these conventional methods, we capitalize on the stochastic nature of the GAN generator.

For instance, DGMs can enhance model pre-training using synthetic data, as well as boost performance in zero-shot and few-shot image classification, as shown in [16]. We amplify the dataset by generating CE multiple times from a single pilot measurement. This approach not only improves the robustness of the dataset but also compensates for occasional inaccuracies in channel estimates from the GAN. Even if the CE are not always perfect, the set of CE can be sufficient for learning BF. By harnessing the stochastic properties of the GAN, we amplify the dataset and enable more effective BF with less information.

We also implement early-stop during the inference epoch. The data amplification discussed earlier increases time consumption, leading us to consider reducing the epoch count during GAN inference. For the convergence in $\mathbf{z}$ in (7), 500 epochs of stochastic gradient descent are required empirically. Therefore, the experiment has two different epoch counts: 500 and 50. While the 500-epoch GAN ensures convergence of (7), the 50-epoch GAN does not fully converge. This early-stop approach offers two main advantages. First, the runtime of the inference loop is reduced by a factor of 10. Second, the approach potentially enhances the robustness of the dataset due to the diversity in CE results.

The baselines of the above method are CS methods, including OMP [8] and EM-GM-AMP [9], which can serve as direct alternatives to GAN for ill-posed inverse problems. Once the pilot signal is received, CS-based CE might produce comparable results, especially for the beamspace sparse simulated channels like DeepMIMO. This suggests that in real-world settings, CS might fall short. On the other hand, GAN-aided CE can potentially excel over CS techniques in BF by amplifying the dataset. Our analysis compares the performance of GAN-BAE against that of OMP-BAE and EM-GM-AMP-




Fig. 5: Average SNR of beamforming vs. Data amplification rate for $P_U = 0$ dB.

BAE. The latter two use CS methods for CE and BAE to learn BF.

The performance metric for beam alignment is the SNR by the optimized BS and UE beams, defined as

$$\text{SNR} = \mathbb{E}_{\mathbf{H} \in \mathcal{H}} \left[10 \log \frac{P_B |\mathbf{w}^H \mathbf{H} \mathbf{f}|^2}{|\mathbf{w}^H \mathbf{n}|^2} \right] \quad (9)$$

IV. RESULTS & DISCUSSIONS

In this section, we highlight the enhanced performance of GAN-BAE over other benchmarks when $N_{\text{pilot}} = 8$. The performance consistently outperforms across different ranges of UE power. Interestingly, even if the NMSE value of the CE is higher, it yields a superior BF gain. This suggests that a low NMSE does not necessarily assure finding the best beamforming direction.

A. Gain of the GAN-BAE with amplified dataset vs baselines

The GAN-BAE performance is measured using the average SNR in dB. The standard setting for our experiments is $P_U = 0$ dB. In Fig. 5a for the $N_{\text{pilot}} = 8$ scenario, GAN-BAE outperforms OMP-BAE by 2.8 dB, and EM-GM-AMP-BAE by 4.7 dB. This is because pilot measurements are not sufficient to fully reconstruct the channel with CS methods, as it is half of the UE channel dimension N_U . On the other




Fig. 6: Average SNR of beamforming vs. UE power P_U for $N_{\text{pilot}} = 8$, Data amplification rate = 1,000.

hand, Fig. 5b with $N_{\text{pilot}} = 16$, the baseline shows similar performance with GAN-BAE. This is attributed to providing enough pilot measurements to capture the full UE channel dimension, N_U .

Fig. 5 illustrates that the performance of the OMP-BAE and the EM-GM-AMP-BAE significantly differs with N_{pilot} . For $N_{\text{pilot}} = 16$, the SNR of the OMP-BAE and the EM-GM-AMP-BAE show a minor gap between the SNR and the BAE trained using actual channels. Yet, with $N_{\text{pilot}} = 8$, both experience a marked decline in SNR. On the other hand, the GAN-BAE shows a high average SNR on both $N_{\text{pilot}} = 8, 16$. This underlines the ability of the GAN-BAE to deliver a more stable, higher SNR when having a lower N_{pilot} . Our analysis will further focus on the case of $N_{\text{pilot}} = 8$ when the number of the UE measurements is insufficient.

As we increase the data amplification rate in Fig. 5, we can observe SNR improves significantly. The data amplification rate of 1,000 boosts the performance by 2.1 dB compared to the data amplification rate of 1. Given the limited pilot measurements, this gap arises not from adding new information but from leveraging the inherent randomness of a DGM, which effectively captures the channel distribution.

Calculating the CE from pilot measurements presents an ill-posed inverse problem, yielding multiple potential solutions. Some of these solutions can significantly deviate from realistic channels. The GAN generator ensures that these estimates align closely with realistic channels, providing ample candidate channel matrices. This allows us to determine the optimal beam angles based on an average-case channel matrix derived from solving the inverse problem.

In Fig. 5a, we now examine GAN inference using 500 iterations which is fully iterated, and 50 iterations which is early stopped. Given the extended time required for data amplification, an approach to compensate for this is to terminate the GAN inference iteration prematurely. This allows for quicker dataset generation and, as an added benefit, yields a more diverse channel generation without drastically violating the channel distribution. An early-stopped CE results in an SNR increase of 0.1 to 0.3 dB compared to a fully converged




Fig. 7: CE NMSE vs. UE power P_U for $N_{\text{pilot}} = 8$. It shows that NMSE is not reflective when the channel is solely used to calculate the single optimal beam.

CE of a ten times lower data amplification rate. This suggests that instead of investing time and resources to fully converge the GAN-based CE, stopping it early and producing more samples for the BAE training is more efficient.

In Fig. 6, the GAN-BAE method consistently outperforms the baseline methods across all UE power levels ranging from -10 dBm to 20 dBm. It shows that the method is not confined to specific SNRs but remains robust across different UE power levels.

B. Is NMSE a good metric for channel estimation?

Although Fig. 5 demonstrates that the GAN-BAE has the highest SNR, GAN does not deliver the “best” CE by conventional standards. Fig. 7 indicates that the OMP-based CE achieves the lowest NMSE value, which is lower than the fully converged GAN-based CE and EM-GM-AMP-based CE by 2 dB and surpasses the early-stopped GAN-based CE by 4 dB – a significant difference. Why does the GAN-BAE outperform OMP-BAE and EM-GM-AMP-BAE?

Relying solely on the average of the NMSE value can be misleading. When reconstructed channels deviate significantly from the expected distribution, their impact tends to be minimized by the averaging process. Given the ill-posed nature of the problem and the limited number of pilot measurements with CS methods, there is a risk that the reconstructed channel could differ significantly from the desired outcome. When averaged, a single deviation can be concealed, presenting an inherent risk when training models on beamforming. Even a single outlier can mislead the system into focusing its energy on an unintended angle. Hence, when the primary objective of CE is to pinpoint a singular optimal beamforming direction, average NMSE might not serve as the most indicative metric.

V. CONCLUSION

We proposed a novel framework for deep learning aided beam alignment, leveraging a combination of GAN-based channel estimator and generator, and a site-specific beam alignment engine (BAE). Importantly, this GAN-BAE framework requires only uplink pilot measurements to train the

entire model, making it considerably closer to “deployment ready” than most prior approaches that require extensive synthetic simulations to produce a dataset for offline training. Furthermore, the heavy computational tasks are at the BS side, reducing the computational burden on the UE and minimizing UL transmissions. Our experimental results show significant improvement in beamforming gain over traditional CS techniques such as the OMP and the EM-GM-AMP. Moreover, the GAN-BAE framework demonstrates robust performance in various power levels of UE.

Potential future directions include handling diverse UE distributions in real time. Is it reasonable to make biased beam predictions based on changed distribution? Determining the optimal retraining interval, especially given the varying pilot measurement presents future challenge topics that are central to the future development of our proposed framework.

REFERENCES

- [1] Y. Heng, J. G. Andrews, J. Mo, V. Va, A. Ali, B. L. Ng, and J. C. Zhang, “Six key challenges for beam management in 5G and 6G systems,” *IEEE Commun. Mag.*, vol. 59, no. 7, pp. 74–79, Jul. 2021.
- [2] M. Giordani, M. Mezzavilla, and M. Zorzi, “Initial access in 5G mmWave cellular networks,” *IEEE Commun. Mag.*, vol. 54, pp. 40–47, Nov. 2016.
- [3] Y. Heng, J. Mo, and J. G. Andrews, “Learning site-specific probing beams for fast mmWave beam alignment,” *IEEE Trans. Wireless Commun.*, vol. 21, no. 8, pp. 5785–5800, Jan. 2022.
- [4] A. Ali, N. González-Prelcic, and R. W. Heath, “Millimeter wave beam-selection using out-of-band spatial information,” *IEEE Trans. Wireless Commun.*, vol. 17, no. 2, pp. 1038–1052, Feb. 2018.
- [5] Y. Zhang, M. Alrabeiah, and A. Alkhateeb, “Reinforcement learning of beam codebooks in millimeter wave and terahertz MIMO systems,” *IEEE Trans. Commun.*, vol. 70, no. 2, pp. 904–919, Nov. 2021.
- [6] Y. Heng and J. G. Andrews, “Grid-free MIMO beam alignment through site-specific deep learning,” *IEEE Trans. Wireless Commun.*, Jun. 2023, [Online early access].
- [7] A. Alkhateeb, “DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications,” in *Proc. Inf. Theory and Appl. Workshop (ITA)*, Feb. 2019, pp. 1–8.
- [8] A. Alkhateeb, O. E. Ayach, G. Leus, and R. W. Heath, “Channel estimation and hybrid precoding for millimeter wave cellular systems,” *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 831–846, Oct. 2014.
- [9] J. P. Vila and P. Schniter, “Expectation-maximization Gaussian-mixture approximate message passing,” *IEEE Trans. on Signal Processing*, vol. 61, no. 19, pp. 4658–4672, Jul. 2013.
- [10] E. Balevi and J. G. Andrews, “One-bit OFDM receivers via deep learning,” *IEEE Trans. Commun.*, vol. 67, no. 6, pp. 4326–4336, Jun. 2019.
- [11] J. Zicheng, G. Shen, L. Nan, P. Zhiwen, and Y. Xiaohu, “Deep learning-based channel estimation for massive-MIMO with mixed-resolution ADCs and low-resolution information utilization,” *IEEE Access*, vol. 9, pp. 54 938–54 950, Apr. 2021.
- [12] A. Doshi, M. Gupta, and J. G. Andrews, “Over-the-air design of GAN training for mmWave MIMO channel estimation,” *IEEE J. Sel. Areas Inf. Theory*, vol. 3, no. 3, pp. 557–573, Sep. 2022.
- [13] M. Arvinte and J. I. Tamir, “MIMO channel estimation using score-based generative models,” *IEEE Trans. Wireless Commun.*, vol. 22, no. 6, pp. 3698–3713, Nov. 2022.
- [14] T. L. Marzetta, “Noncooperative cellular wireless with unlimited numbers of base station antennas,” *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 3590–3600, Oct. 2010.
- [15] A. Bora, E. Price, and A. G. Dimakis, “AmbientGAN: Generative models from lossy measurements,” in *Proc. Int. Conf. Learn. Representations*, Feb. 2018, pp. 1–22.
- [16] R. He, S. Sun, X. Yu, C. Xue, W. Zhang, P. Torr, S. Bai, and X. Qi, “Is synthetic data from generative models ready for image recognition?” in *Proc. Int. Conf. Learn. Representations*, Feb. 2023, pp. 1–24.