

Machine Learning Codebook Design for Initial Access and CSI Type-II Feedback in Sub-6GHz 5G NR

Ryan M. Dreifuerst, *Student Member, IEEE*, and Robert W. Heath Jr. *Fellow, IEEE*

Abstract—Beam codebooks are a recent feature to enable high dimension multiple-input multiple-output in 5G. Codebooks comprised of customizable beamforming weights can be used to transmit reference signals and aid the channel state information (CSI) acquisition process. Codebooks are also used for quantizing feedback following CSI measurement. In this paper, we unify the beam management stages—codebook design, beam sweeping, feedback, and data transmission—to characterize the impact of codebooks throughout the process. We then design a neural network to find codebooks that improve the overall system performance. The proposed neural network is built on translating codebook and feedback knowledge into a consistent beamspace basis similar to a virtual channel model to generate initial access codebooks. This beamspace codebook algorithm is designed to directly integrate with current 5G beam management standards without changing the feedback format or requiring additional side information. Our simulations show that the neural network codebooks improve over traditional codebooks, even in dispersive sub-6GHz environments. We further use our framework to evaluate CSI feedback formats with regard to multi-user spectral efficiency. Our results suggest that optimizing codebook performance can provide valuable performance improvements, but optimizing the feedback configuration is also important in sub-6GHz bands.

I. INTRODUCTION

Fifth-generation (5G) cellular systems have adopted a beam-based approach for multiple-input multiple-output (MIMO) communications. In such systems, broadcast control signals are beamformed to provide array gain [1]. This helps increase coverage and meet link budgets at high carrier frequencies. While the motivation for beam-based systems was to provide array gain in millimeter-wave (mmWave) frequency bands, the same framework is also used at low frequencies. Unfortunately, optimization of the beamforming framework, and subsequent feedback strategies, has not been extensively investigated for sub-6GHz multi-user MIMO (MU-MIMO), despite

the importance of low- and mid-band MIMO in current and anticipated future standards [2].

Codebook-based beam management is a framework for beamforming and feedback that enables transmitting beamformed reference signals and multiple forms of CSI acquisition [1]. Unprecoded reference signals were used up through 4G LTE Release 14 because multi-antenna beamforming was not needed to achieve higher signal-to-noise ratio (SNR) for synchronization and channel estimation [3]. Beamformed control signals have become a necessity in mmWave bands as a result of the increasing bandwidths (which reduce the SNR with constant transmit power) and shrinking antennas sizes—resulting in pre-beamformed SNR on the order of -15dB or less [4]. By beamforming the control and reference signals, receivers obtain more accurate synchronization and channel measurements.

The process of obtaining CSI at the transmitter involves multiple steps and codebooks to balance the overhead, latency, and accuracy of the CSI [5], [6]. To begin the process, the base station (BS) transmits a synchronization signal block (SSB) with a beamformer selected from a codebook (SSB codebook) that enables synchronization and random access. The SSB codebook is typically small and all of the codewords in the codebook are used during each SSB period. The user equipment (UE) will provide a small amount of feedback to conclude the initial access that includes a beam selection from the SSB codebook. The BS will use this information to select a subset of beamformers from a large codebook, different from the SSB codebook, for transmitting channel state information reference signals (CSI-RS). The relationship between these two codebooks can be understood from the perspective of a hierarchical beam search [7]. The CSI-RS enable hybrid beam training by providing multiple precise beams for analog beam training—often referred to as beam refinement—and pilots for digital channel estimation. The UE will provide an in-depth packet of feedback following CSI-RS transmission that includes the CSI-RS codebook beamformer selection as well as information based on the channel estimation. While the first iterations of beam management used wide and nar-

Ryan M. Dreifuerst and Robert W. Heath Jr. are with North Carolina State University, Raleigh, NC 27695 {rmdreifu, rwheathjr}@ncsu.edu. This material is based upon work supported by the National Science Foundation under grant nos. NSF-ECCS-2153698, NSF-CCF-2225555, NSF-CNS-2147955 and is supported in part by funds from federal agency and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program.

row discrete Fourier transform (DFT) beamformers [4], there is no requirement that the beamformers be “wide” and “narrow” for the SSB and CSI-RS codebooks. In fact, the codebooks can be arbitrary since the UEs do not need knowledge of the beamformers during initial access, beam refinement, or even data transmission.

While the SSB and CSI-RS codebooks can be arbitrary, the codebook used for quantizing the feedback (FB codebook) is precisely defined depending on the format (type-I, type-II). The FB codebook is specific for BS and UE array geometries so that the CSI can be represented and fed back as indices corresponding to a known codebook at the BS and UE. It is not strictly necessary that even the feedback codebook is used when reciprocity exists, as reciprocity between the uplink and downlink channel can be exploited in time-division-duplexing systems. Feedback from downlink training is still beneficial in these systems, however, due to asynchronous link budgets [8] or when the number of transmitting and receiving antennas on a UE is not equal [9].

In our work, we unify the two stages of codebook design with the feedback framework and MU-MIMO data transmission to achieve better network performance. To demonstrate the applicability of our work we present a system-level evaluation of the beam management framework in sub-6GHz bands. Our contributions in this manuscript are summarized as follows

- First, we precisely describe the beam management and CSI type-II feedback approach for 5G networks. This is the first work that unifies the beam management stages—codebook design, beam training, feedback, and data transmission—into a joint, MU-MIMO system model. We explain how the various codebooks are used, what is specified in the standards versus left up to implementations, and provide an efficient implementation for feedback quantization in CSI type-II.
- Second, we present a novel neural network architecture, BeamSpace-Codex (BSC), and feedback processing technique for SSB codebook generation. We constrain the system to follow 5G beam management timing which presents a significant constraint for dynamic SSB and CSI-RS codebook generation. To mitigate this constraint, the SSB codebook is dynamically generated—which is manageable at the SSB periodicity—and the CSI-RS codebook is determined as a precomputed DFT decomposition of the SSB codebook. With this integration, we join the arbitrary codebooks learned via deep learning with the feedback quantization technique to facilitate a hierarchical beam search that improves the beam management process used in 5G new radio (NR).
- Third, we present the results of an extensive eval-

uation of our proposed method with benchmark DFT codebooks and CSI-based eigen-beamforming solutions using various levels of feedback and overhead. We show that our proposed solution achieves significant performance improvements, extends to different user numbers, and does not need to be relearned for new array sizes. We also show that MU-MIMO performance is dramatically limited by the quantization resolution techniques in low-band 5G. In contrast, the frequency selectivity and resource allocation have little impact in the current framework and scattering environment.

There is significant work on beam training and beam alignment. Here we review a set of relevant prior work relating to machine learning and beam training. In [10], a gradient-optimization approach, not using deep learning, is shown to be successful in aiding the beam search process with respect to maximizing the SNR in 60GHz personal networks. Another approach using classical data-driven techniques is presented in [11] for vehicular networks based on historical beam training results. While the codebooks and algorithms are relatively straightforward, the idea of “learning” an efficient codebook from site-specific data is a cornerstone of many modern beam alignment techniques [12]–[14]. These papers [10]–[14] all focus on mmWave channels and often exploit the sparsity of such environments to aid the beam training process. Furthermore, all of these papers focus exclusively on beam alignment with standard orthogonal codebooks like DFT and phase-shifted codebooks. While DFT codebooks are sensible for certain array geometries and limited scattering environments at mmWave bands, these codebooks do not necessarily work well in rich scattering. Furthermore, none of these papers consider a MIMO format with multiple UE antennas and streams of data. In this paper, we consider a full MU-MIMO OFDM system in a rich scattering environment and use dynamic codebooks.

In another line of related work, algorithms have been proposed that design beamformers or codebooks [7], [13], [15]–[20]. Hierarchical codebooks and the associated design process were proposed in [7], [15], [16] which separated the beam training into multiple stages of wide-beam and narrow-beam codebooks. A similar strategy was adopted by 5G using the SSB and CSI-RS codebooks and plays an important role in systems designed for 5G deployment. Deep learning was applied to the task in [17], which uses deep reinforcement learning to design broadcast beam patterns. While designed for small-scale MIMO, the results suggest that arbitrary learned beam patterns can be an effective way to initiate communication in broadband networks. A supervised and an unsupervised learning approach were presented in

[18], which showed promising results when used to design beamforming vectors from channel measurements. These results were based on a simplified channel model, though, and perfect CSI was assumed at the BS. Recently, a deep, federated, reinforcement learning strategy was proposed [13] that jointly trained a beam management model over a network using user location information. The focus was on deciding which sectors would be active, thereby narrowing the “codebook,” although the paper ignores all beamforming besides sector-level association and large-scale pathloss. Another paper considered the task of learning probing beams [19], which is one of the only papers to consider both the SSB and CSI-RS codebooks. The results show significant gains can be achieved over a basic hierarchical search, although perfect CSI is assumed to be available for training and the channels are assumed to be narrowband. The results from [19] provide motivation for researching realistic deep learning codebook methods. In our previous work we designed a neural network for SSB codebook generation in narrowband channels as well but trained using an angular representation of the channel [20]. The results, however, did not extend well to wideband systems and did not consider the impact that the SSB codebook has on system-level performance. In this investigation, we now model wideband channels and fully incorporate the entire beam management and data transmission into our evaluation.

An important part of the 5G beam management framework is the connection between beam training and feedback. An enhanced CSI feedback strategy was proposed for full-dimension MIMO [3] in 4G Release 13 which is based on a linear combination of DFT codewords. This strategy is very similar to the type-II feedback format introduced in 5G Release 15. One key difference is that the initial idea assumed the users must know the beamforming vector used during pilot transmission, which was not adopted into the standards. CSI type-I, which is primarily for SU-MIMO, was investigated with UE location [21], and as an indicator for beamforming vector selection [22]. These works are focused on single-user data transmission, however, and do not consider multi-stream precoding. There has also been growing interest in using machine learning to design or modify the feedback [23]–[25]. All those works [23]–[25] assume that both the network and the users can share models for encoding and decoding the feedback, which is not supported or easily introduced into network operation. In this paper, we do not use deep learning to change the feedback methods. Instead, we characterize the relationship and limitations the feedback and beam management strategy have on MU-MIMO performance.

Notation: $\mathbf{A}$ is a matrix, $\mathbf{a}$ and $\{a[i]\}_{i=1}^N$ are column vectors and a, A denote scalars. $\mathbf{A}^T$, $\mathbf{A}^*$, and $\mathbf{A}^\dagger$

represent the transpose, conjugate, conjugate transpose, and pseudo-inverse of $\mathbf{A}$. $\mathbf{A}[k, \ell]$ denotes the entry of $\mathbf{A}$ in the k^{th} row and the ℓ^{th} column. The same meaning is also associated with $\mathbf{A}_{k, \ell}$. Similarly, $\mathbf{A}[:, k]$ refers to the k^{th} column of $\mathbf{A}$. Unspecified norm equations are $\|\mathbf{a}\|_2 = \sqrt{\mathbf{a}^* \mathbf{a}}$ for vectors and the Frobenius norm $\|\mathbf{A}\|_F = \sqrt{\text{Tr}(\mathbf{A} \mathbf{A}^*)}$ for matrices. Superscripts with text font are used to describe a variable, while italic font is used for exponentiation. We make use of the disjoint operator $\mathbf{A} \setminus \alpha$ to describe the matrix (or vector) $\mathbf{A}$ excluding the subset (or element) α . The operator $\mathbb{E}[\cdot]$ is used for the expectation of a random variable. Due to the notational complexity of MU-MIMO with OFDM, we will always use u to refer to a specific UE, t as a specific time, k as a specific frequency resource, and n_t/n_r to refer to a specific transmit or receive antenna.

The remainder of our paper is organized as follows. First, we introduce the system model beginning from an arbitrary MU-MIMO OFDM channel and continuing through the SSB and CSI-RS beam management processes. We conclude this section with a description of the channel estimation, propose a fast feedback determination algorithm, and define the data transmission stage. Next, we introduce the beamspace observation and neural network architecture that form the second contribution of our work. In the final sections, we present the simulation setup and evaluate the various codebooks and feedback formats to address the lack of prior work when it comes to codebooks and CSI formats in sub-6GHz 5G NR.

II. SYSTEM MODEL

We begin the investigation with an overview of the channel definition and assumptions before introducing the beam management steps including SSB transmission, CSI-RS transmission, feedback, and downlink data transmission, as shown in Figure 1. Throughout this paper, we will limit the system configuration to a single cell with multiple UEs each equipped with multiple antennas and all arrays are fully digital and single-polarization. While the inclusion of multiple cells is more realistic, the impact of multi-cell interference on beam training and feedback can be mitigated with limited coordination between nearby cells, although the interference is underestimated as a result of the lack of interfering BS. The use of single-polarized arrays is also a simplification, although during beam training the second polarization is not significantly different for the system operation. We will investigate multiple BS deployments with dual-polarized arrays in future work.

A. MU-MIMO channel model

We model the system so that it is representative of real-world conditions to evaluate performance in a realistic

least the same size, where N is typically large but only a small subset of beamformers is selected to be used in one period. The FB codebook $\mathbf{B}^{\text{DFT}} \in \mathbb{C}^{N_P \times N_X O_H N_Y O_V}$ is an oversampled DFT codebook with horizontal and vertical oversampling factors O_H and O_V . Previously, 4G used a feedback codebook based upon predefined matrices called CSI type-I precoding matrix indicator (PMI). 5G augmented the feedback to include type-II PMI which corresponds to a set of directional array responses with amplitude and phase combining factors. Each of these codebooks is designed to serve different purposes, although there are many inter-relationships that make designing and evaluating codebooks challenging.

Designing codebooks for 5G is challenging because each codebook serves a complementary role with the others for obtaining CSI. Codebook design in the physical layer is often incorrectly treated as a one-dimensional problem focused on maximizing the received power. Beamforming codebooks affect subsequent system operations that impact the overall data transmission. For example, it may be assumed that the beamforming, feedback, and data transmission all use the same codebook, but such a simplification is not required or realistic. We consider such a scenario in our simulation results in Figure 5(b), but the performance is significantly worse than traditional, complementary codebooks. This motivates bringing together the various codebooks and beam management into a multi-step system model.

C. SSB transmission and reception

Initial access is used in mobile networks to obtain limited channel information, achieve synchronization, and set up random access procedures. In 5G NR, a cell may initiate the initial access period at regular intervals of $\{5, 10, 20, 40, 80, 160\}$ ms [6], to control the periodicity a UE must be active (not in power-saving mode) and provide feedback to the network. During this period, the cell will transmit SSBs that contain primary and secondary synchronization signals (PSS, SSS) as well as demodulation reference signals (DMRS) [6]. These SSBs are beamformed using a specific beamforming vector selected from $\mathbf{B}^{\text{SSB}}$ and associated with a beam index. Each SSB is transmitted sequentially in time, using 20 contiguous resource blocks. This is a much smaller bandwidth than the downlink data transmission and is a key limitation of relying on SSB feedback to obtain CSI in frequency-selective channels. Furthermore, the resource allocation of SSBs is essentially fixed to assist new users with joining a network. Depending on the cell carrier frequency and subcarrier spacing, a cell may transmit up to $L_{\text{max}} = \{1, 4, 8, 64\}$ beams in a period, and all cells must transmit at least one beam. During transmission, the UE will measure the Reference Signal Received Power (RSRP) and report the measurement using the random

access channel slot corresponding to the index of the beam with the highest RSRP. This way the BS knows how strong the signal was and to which of the SSB codebook's codewords it corresponds.

The SSB is transmitted over a band of $\mathbf{K}_{\text{SSB}}$ frequency resources at times $\mathbf{T}_{i,\text{SSB}}$ for the i^{th} SSB using beamformer $\mathbf{f}_i \in \mathbf{B}^{\text{SSB}}$. While transmitting the reference signal the rest of the band is still available for data transmission, so the system power is shared over all K resource blocks. In total, all L_{max} SSBs are transmitted from the codebook $\mathbf{B}^{\text{SSB}}$. During SSB reception, the generic received signal model from (2) is modified to account for the single data stream

$$y_{u,t,k}^{\text{SSB}_i} = \frac{1}{\sqrt{KN_T}} \mathbf{w}_{u,t,k}^* (\mathbf{H}_{u,t,k} \mathbf{f}_i s_{t,k} + \mathbf{N}_{u,t,k}). \quad (3)$$

In mmWave bands, it is common that $\mathbf{w}$ is also determined through beam training. In sub-6GHz bands, the receiver instead uses antenna selection or simply the first antenna for power saving [27]. With the received signal model (3) in mind, we now define an important quantity provided as feedback during initial access, the RSRP.

D. Reference signal receive power

RSRP is one of the primary metrics that the receiver will measure during initial access for determining the channel quality. The BS will use the RSRP feedback to determine the strength of the signal, the initial code rate to be used, and what CSI-RS beamformers should be employed at the next iteration. The RSRP measures the received power during a given reference signal (SSB or CSI-RS). In the SSB case, the DMRS, known at the transmitter and receiver, is a pilot signal used for RSRP measurement and decoding the master information block provided in the SSB. We will assume antenna selection for the $n_r \in \{0, \dots, N_R - 1\}$ antenna during SSB reception, but because it is fully digital the UE can determine the antenna selection after receiving the full SSB and selecting the strongest receive branch. The RSRP is then

$$\mathbf{y}_{u,t,k}^{\text{SSB}_i} = \frac{1}{\sqrt{KN_T}} \mathbf{H}_{u,t,k} \mathbf{f}_i s_{t,k} + \mathbf{N}_{u,t,k} \quad (4)$$

$$\text{RSRP}_{i,u} = \max_{n_r} \sum_{k \in \mathbf{K}_{\text{DMRS}}} \sum_{t \in \mathbf{T}_{i,\text{DMRS}}} \left\| \mathbf{y}_{u,t,k}^{\text{SSB}_i} \right\|^2. \quad (5)$$

Note that (5) is for a specific UE (u) and SSB (i), so the total SSB information aggregated at the BS is comprised of the strongest RSRP and corresponding SSB index (SSBRI) for each of the UEs, denoted as $(\mathbf{p}, \mathbf{m})$, and defined as

$$\mathbf{p} = \left\{ \max_i \text{RSRP}_{i,u} \right\}_{u=0}^{U-1} \quad (6)$$

$$\mathbf{m} = \left\{ \arg\max_i \text{RSRP}_{i,u} \right\}_{u=0}^{U-1}. \quad (7)$$

It is worth noting that the actual SSBRI is not transmitted by the users; the random access slot that the UE responds during will correspond to a specific SSBRI. This is an efficient way to reduce the amount of feedback, allow the BS to perform receive combining for the UE feedback, and still share the necessary information.

E. CSI-RS transmission and reception

Beam refinement is the second step of the hierarchical beam management process that is used to obtain better beam alignment and enable wideband channel estimation. Beam refinement is done through the transmission of CSI-RS which is an encompassing term used for the reference pilots, the beamformed reference signal, and the logical ports that the pilots are sent on. In contrast to SSB signals, CSI-RS are extremely flexible with many combinations of time and frequency resources allocatable. There are some limitations on how often CSI-RS can be transmitted, but the default periodicity a UE expects when joining the network is 80ms and the symbols must span a larger bandwidth than the SSB—at least 24 resource blocks and up to the entire bandwidth should contain CSI-RS symbols for a measurement report. Furthermore, CSI-RS are also transmitted both periodically and aperiodically to assist beam tracking with highly mobile users. All told, the beam refinement stage requires more resources than initial access, which is why a hierarchical beam search is used to limit the number of CSI-RS processes needed to achieve a strong connection.

Once a BS receives the SSB feedback for a new user, it can attempt to gather more precise CSI through channel state information reference signals. The CSI-RS (and associated CSI-RS beamformers $\mathbf{F}^{\text{CSI-RS}}$) are used for refined channel estimation, although the estimation is performed on the beamformed channel $\mathbf{H}_u \mathbf{F}^{\text{CSI-RS}}$, rather than the physical channel. This enables larger hybrid arrays and improves the SNR for more accurate estimation. Readers are encouraged to review [5], [6], [9] for more details of CSI-RS. The BS can use the SSB feedback to determine and transmit a set of CSI-RS signals with beamformers that are intended to further increase the RSRP compared to the SSB.

The design of CSI-RS codebooks has been studied to some degree, but DFT beams are common [3], [13], [28] (and references therein). Often, it is assumed that the set of beamformers is selected from the same codebook as the codebook used for feedback, which is specified as an oversampled DFT codebook. There is no specification or need for the SSB or CSI-RS codebooks to be DFT codebooks, although using DFT beamformers during CSI-RS has advantages related to fast and efficient processing. For example, the CSI-RS beamformers can be efficiently

selected as the largest DFT components of the SSB beamformers. We will assume the CSI-RS codebook, $\mathbf{B}^{\text{CSI-RS}}$, is comprised of $N_X N_Y O_H O_V$ oversampled DFT vectors throughout this paper, although additional investigations into joint SSB–CSI-RS codebook algorithms are planned in future work. The selection of the CSI-RS beams from SSB decomposition is further refined in Section III.

The CSI-RS process involves transmitting a set of pilots $\mathbf{s}$ on the N_P ports along with P_{CSI} beamformers $\mathbf{F}^{\text{CSI-RS}} = \{\mathbf{f}_i\}_{i=0}^{P_{\text{CSI}}} \subseteq \mathbf{B}^{\text{CSI-RS}}$ chosen by the serving BS to improve the UE channel estimation SNR for the effective channel. The beamformers are used during non-overlapping time-frequency resources ($\mathbf{T}^{\text{CSI}_i}, \mathbf{K}^{\text{CSI}_i}$), with the aggregate set of beamformers corresponding to what we define as $\mathbf{F}^{\text{CSI-RS}}$. Each CSI-RS allocation will have at least P_{CSI} OFDM resource elements allocated for channel estimation in one resource block and only one CSI-RS beam is active for a given resource to allow for a similar, interference-free beam selection as the SSB process. This is done by setting all other CSI-RS resources and ports to be zero-power CSI-RS except the active one for the specified resources. In a dual-polarized system, the two polarizations are transmitted together to estimate a co-polarization factor that would be fed back during CSI reporting. The whole CSI-RS process is allocated over NRB resource blocks, which is a configurable parameter to use wider bandwidths and noise averaging during channel estimation. The beam reception and selection process is identical to the RSRP and SSB selection in the previous subsection. For more information on the CSI-RS beam and port allocation please see [3] and references therein. At this stage, the entire UE antenna array is active and the UE can perform maximum ratio combining $\mathbf{w}_{u,t,k}^i = (\widehat{\mathbf{H}_{u,t,k}} \mathbf{f}_i)^*$ based on the estimated beamformed channel and appropriately normalized. The estimated SNR at the receiver is

$$\text{SNR}_u = \max_{c \in P_{\text{CSI}}} \frac{1}{K} \sum_{k \in \mathbf{K}_{\text{CSI}_c}} \sum_{t \in \mathbf{T}_{\text{CSI}_c}} \frac{1}{N_T} \frac{\|\widehat{\mathbf{H}_{u,t,k}} \mathbf{f}_c\|^2}{\mathbb{E}[\|\mathbf{N}_{u,t,k}\|^2]}. \quad (8)$$

Note that the noise power $\mathbb{E}[\|\mathbf{N}_{u,t,k}\|^2]$ can be estimated using zero-power CSI-RS resources. The SNR can be reported during feedback, although one of the most important aspects of the SNR is how the channel estimation (and subsequent feedback) is impacted. For the next part of this section, we will focus on a single user u with SNR corresponding to the channel estimation SNR for the beamformed channel ($\mathbf{H}_{u,t,k} \mathbf{F}^{\text{CSI-RS}}$).

F. Channel estimation

Channel estimation accuracy depends on the number of pilot symbols and the received SNR of the pilots.

The algorithm for UE channel estimation is not defined in the specification and is not critical for our system-level analysis. For simplicity, we will assume channel estimation via a simple least squares algorithm so that only knowledge of the pilots is required. Now, we motivate the inclusion of channel estimation in beam management by considering the relationship between the channel estimation error, the codebook-dependent-SNR, and the number of pilot resources N_{pilots} . The resulting mean squared error (MSE) of the channel estimate is [29, section 3.7]

$$\text{MSE} = \frac{1}{\frac{N_{\text{pilots}}}{N_T} \text{SNR}}. \quad (9)$$

Therefore, we can see that a primary feature of the beam management process is to enable more accurate channel estimation by beamforming the reference pilot signals. It might seem that the overhead of using multiple beams is wasted given that the number of pilots could alternatively be increased according to (9). The issue with such logic is that: 1) it assumes that the signal is always detectable over the noise floor, which is not necessarily the case for non-beamformed signals and 2) the beamforming can provide a gain on the order of 20dB or more, yet that would require 100 times more pilot symbols, far larger than the number of beams used.

Calculating the estimated effective channel, $\widehat{\mathbf{H}\mathbf{F}}$, which may not be frequency selective, is the first step in determining the precoder matrix information (PMI), which is one of the quantized CSI elements provided during feedback. The PMI can correspond to different aspects of the estimated channel because it is left to operator implementation as part of the PMI selection procedure. For example, the user could provide PMI according to its desired beamformers (calculated via singular value decomposition of $\widehat{\mathbf{H}\mathbf{F}}$). The BS could then use this information to beamform to a given user and null steer [30] the beams for other users. This puts the majority of the computational burden, however, on the UE. Instead, we follow models similar to [31] where the user provides PMI corresponding to the effective channel estimate, rather than the SVD of the effective channel estimate. This way the BS handles most of the computation, channel reconstruction, and precoder determination.

G. CSI feedback and reconstruction

We will focus on CSI type-II feedback (Release 16), as our feedback baseline, which is intended for MU-MIMO with up to 2 layers per user [27], and is also compliant with Release 17. Although not presented here, codebooks can also be used for SU-MIMO feedback; we use this form of feedback when comparing against the SU-MIMO results in Section IV. The UE quantizes the channel estimate according to the FB codebook for a

specified L_{CSI} beams per rank up to the rank indicator, which is the maximum supported MIMO rank, R . A straightforward method for selecting the PMI would be to exhaustively try every possible combination of L_{CSI} beams from the codebook and select the best one (according to some metric) as typically done in open source packages [32]. This is a significant computational burden on UEs for large L_{CSI} due to the combinatorial growth of the search space. We propose using the orthogonality of DFT codebooks to reduce the computational complexity compared to exhaustively searching.

First, we define the oversampled DFT codebook $\mathbf{B}^{\text{DFT}}$ beginning with the (non-oversampled) DFT codebook for a uniform linear array $\mathbf{U}_N$ defined in the same way as e.g. [33]. Then we can define a diagonal oversampling matrix $\mathbf{D}_{O,N}$ for an oversampling O and N elements as

$$\mathbf{D}_{O,N} = \text{diag} \left(1, e^{(-\frac{j2\pi}{ON})}, \dots, e^{(-\frac{j2\pi(N-1)}{ON})} \right). \quad (10)$$

Then, the oversampled DFT codebook $\mathbf{B}_{O,N} \in \mathbb{C}^{N \times ON}$ for a one-dimensional array is constructed using matrix powers of $\mathbf{D}_{O,N}$

$$\mathbf{B}_{O,N} = [\mathbf{D}_{O,N}^0 \mathbf{U}_N, \mathbf{D}_{O,N} \mathbf{U}_N, \dots, \mathbf{D}_{O,N}^{O-1} \mathbf{U}_N]. \quad (11)$$

Finally, the UPA oversampling codebook can be defined as the kronecker product of the N_X and N_Y dimensions

$$\mathbf{B}^{\text{DFT}} = \mathbf{B}_{O_H, N_X} \otimes \mathbf{B}_{O_V, N_Y} \in \mathbb{C}^{N_X N_Y \times O_H N_X O_V N_Y}. \quad (12)$$

With the FB codebook defined in (12), the PMI selection procedure can begin. The UE first calculates the inner product for all beams from the oversampled codebook $\mathbf{B}^{\text{DFT}}$ and selects the first beam (and by association the oversampling basis) as

$$\mathbf{C} = \widehat{\mathbf{H}\mathbf{F}} \mathbf{B}^{\text{DFT}} \in \mathbb{C}^{N_R \times N_X O_H N_Y O_V} \quad (13)$$

$$Q_0 = \text{argmax} |\mathbf{C}| \quad (14)$$

where the argmax is determined from the absolute value of all elements of $\mathbf{C}$ and the resulting Q_0 is an index corresponding to oversampling powers $Q_{0,X}$ and $Q_{0,Y}$ in the N_X and N_Y directions. Because (non-oversampled) DFT beams are orthogonal, the PMI can be calculated iteratively by selecting the first beam from the entire oversampled codebook $\mathbf{B}^{\text{DFT}}$, reducing the codebook to just the orthogonal subset, and selecting the beams from the subset. The orthogonal subset can be understood as the kronecker product of the oversampling-shifted DFT codebooks

$$\mathbf{B}_{\text{ortho}, Q_0} = \mathbf{D}_{O_H, N_X}^{Q_{0,X}} \mathbf{U}_{N_X} \otimes \mathbf{D}_{O_V, N_Y}^{Q_{0,Y}} \mathbf{U}_{N_Y} \quad (15)$$

With the orthogonal subset, the next step is to select the L_{CSI} beam indices $\mathbf{Q} \in \mathbb{Z}_+^{N_R \times L_{\text{CSI}}}$ and combining coefficients $\mathbf{A} \in \mathbb{C}^{N_R \times L_{\text{CSI}}}$

$$\mathbf{c}_{\text{ortho}} = \widehat{\mathbf{H}\mathbf{F}}_{n_r} \mathbf{B}_{\text{ortho}, Q_0} \in \mathbb{C}^{N_X N_Y} \quad (16)$$

$$\mathbf{Q}_{n_r} = (L_{\text{CSI}})\text{-argmax} |\mathbf{c}_{\text{ortho}}| \quad \forall n_r \quad (17)$$

$$\mathbf{A}_{n_r} = \mathbf{c}_{\text{ortho}} [\mathbf{Q}_{n_r}] \quad \forall n_r \quad (18)$$

$$\mathbf{A}_{n_r} = \frac{\mathbf{A}_{n_r}}{\mathbf{A}[n_r, \text{argmax} |\mathbf{A}_{n_r}|]} \quad \forall n_r. \quad (19)$$

First, the orthogonal beams are “evaluated” in (16), then we make use of the k -argmax operation defined as the k indices for the k largest values in (17) to select the L_{CSI} beam indices. Finally, the combining coefficients are calculated as the complex values of the beam products (18) and normalized according to the largest magnitude coefficient (19). The oversampling basis $Q_{0,X}, Q_{0,Y}$ is reported back in the PMI so that the overall beam selections can be reconstructed at the BS.

The CSI report contains the CSI-RS indicator (CRI) which is the strong CSI-RS beamforming codeword, the rank of the PMI, the RSRP or SNR, and the PMI feedback beam indices $\mathbf{Q}$, as well as 4 bit amplitude and 8-PSK phase values. The amplitude and phase values come from $\mathbf{A}$ which are the corresponding complex values for each $[n_r, \ell], \ell \in L_{\text{CSI}}$ feedback beam. There is an additional overhead reduction step in the specification where instead of indexing each element of $\mathbf{Q}$ from the entire codebook, the set of indices is translated to one of the $\binom{N_X N_Y}{L_{\text{CSI}}}$ combinations and one index corresponding to the oversampling factors of the strongest index ($Q_{0,X}, Q_{0,Y}$) is used. In this paper, we will ignore the amplitude and phase quantization (assuming perfect amplitude and phase knowledge) to focus on the effects of the codebooks, SNR, and feedback. The CSI report can also contain both wideband and narrowband components for a degree of frequency selective precoding. The number of narrowband components (termed subbands in 3GPP nomenclature) is small, at most $S_B = 8$ in FR1, to mitigate the growing feedback and complexity. All of the previous and subsequent steps are defined for $S_B = 1$, but extending to frequency selective feedback and precoding is straightforward. Release 16 also supports enhanced type-II codebooks, which slightly modify the number of bits for various PMI fields as well as adding frequency domain compression via inverse DFT basis vectors. We investigate how the amount of feedback (L_{CSI}, S_B) impacts performance in Section IV.

H. MU-MIMO data transmission

After the CSI-RS period, the BS must determine the precoder that serves a group of users, often with the goal of maximizing the sum rate or proportional fair rate. Until this stage, all of the downlink transmissions have been multi-cast in the sense that the same pilots and data are intended for all users. Therefore, there is no interference assuming a well-designed OFDM cyclic prefix and subcarrier spacing is used with accurate synchronization.

In contrast, the downlink data transmission carries different data for each user, resulting in interference between data streams. The remainder of this section outlines the channel reconstruction process, precoder selection, and achievable spectral efficiency in the downlink channel.

First, because the feedback is type-II, the BS can reconstruct the PMI (corresponding to the beamformed channel estimate here) with the corresponding complex scaling $\mathbf{A}$ as

$$\widehat{\mathbf{H}}\mathbf{F}[n_r] = \sum_{\ell=0}^{L_{\text{CSI}}-1} \mathbf{A}[n_r, \ell] \mathbf{B}^{\text{DFT}}[:, \mathbf{Q}_{n_r}, \ell] \quad \forall n_r. \quad (20)$$

Then, assuming the beamformers were well-chosen such that $\mathbf{F}^{\text{CSI-RS}}$ is right-invertible, channel reconstruction can be performed successfully. We reintroduce the UE u notation here to allow for aggregating all user channels in subsequent steps by

$$\widehat{\mathbf{H}}_u = \widehat{\mathbf{H}}_u \mathbf{F}^{\text{CSI-RS}} \quad (21)$$

In the case of DFT beams, this process simply becomes an IDFT operation [3], which is significantly more efficient than the matrix inverse of an $N_P \times N_P$ matrix. The efficiency at this step is critical to quickly using the feedback to serve users and is one reason we will tailor our algorithm to integrate with a DFT codebook for CSI-RS. Aggregating the results for each user, the BS obtains the downlink channel estimate $\widehat{\mathbf{H}} \in \mathbb{C}^{U \times N_R \times N_P}$

$$\widehat{\mathbf{H}} = \begin{bmatrix} \widehat{\mathbf{H}}_0, \widehat{\mathbf{H}}_1, \dots, \widehat{\mathbf{H}}_{U-1} \end{bmatrix}. \quad (22)$$

The BS can design the precoders for the set of estimated channels using any form of precoding. To reasonably evaluate the system we will assume a regularized zero forcing (RZF) precoder is used with regularization determined to minimize the signal-to-leakage noise ratio [29, section 9.9]. The RZF precoder for each user is

$$\mathbf{F}_u = \frac{(\sum_{i=0}^{U-1} \widehat{\mathbf{H}}_i^* \widehat{\mathbf{H}}_i + U N_T \mathbb{E}[\mathbf{N}_{u,t,k}^2])^{-1} \widehat{\mathbf{H}}_u^*}{\left\| (\sum_{i=0}^{U-1} \widehat{\mathbf{H}}_i^* \widehat{\mathbf{H}}_i + U N_T \mathbb{E}[\mathbf{N}_{u,t,k}^2])^{-1} \widehat{\mathbf{H}}_u^* \right\|}. \quad (23)$$

It is assumed all precoders are appropriately normalized to $\sqrt{\frac{N_T}{U}}$ to maintain the same total power. The noise factor, $\mathbb{E}[\mathbf{N}_{t,k,u}^2]$ can be estimated based on the RSRP/SNR reported and the bandwidth of the downlink signal. The UE will also perform a combining strategy based on the previously determined rank. In this work, we will assume an LMMSE receiver is used to maximize the signal-to-interference noise ratio (SINR) for any precoder, using the embedded DMRS to determine the combined channel and precoder. The resulting SINR at resource element (t, k) for a given user u and layer r with equal power

allocation per user is obtained as

$$\mathbf{I}_{u,t,k} = \left(\sum_{i=0}^{U-1} \mathbf{H}_{u,t,k} \mathbf{F}_i \mathbf{F}_i^* \mathbf{H}_{u,t,k}^* + U N_T \mathbb{E}[\mathbf{N}_{u,t,k}^2] \right)^{-1} \quad (24)$$

$$\text{SINR}_{u,t,k,r} = \frac{1}{K} \frac{\mathbf{F}_u^*[:,r] \mathbf{H}_{u,t,k}^* \mathbf{I}_{u,t,k} \mathbf{H}_{u,t,k} \mathbf{F}_u[:,r]}{1 - \mathbf{F}_u^*[:,r] \mathbf{H}_{u,t,k}^* \mathbf{I}_{u,t,k} \mathbf{H}_{u,t,k} \mathbf{F}_u[:,r]}. \quad (25)$$

While the SINR expression (25) appears complicated, it is a simplified ratio between the signal power for the data stream (u, r) versus the interference power of all other streams. The SINR at this step is not calculated in a physical system, but it is necessary for evaluating the achievable spectral efficiency.

Finally, the most critical metric in the wireless network is the sum SE or sum rate when applied to a specific bandwidth. A fairness constraint can also be applied, however, we have not specified any form of scheduling so we are most interested in maximizing the sum spectral efficiency. Assuming Gaussian signaling and treating interference as Gaussian noise, the achievable spectral efficiency, $\text{SE}_{u,t}$, is

$$\text{SE}_{u,t} = \sum_{k=0}^{K-1} \sum_{r=0}^{R-1} \log_2(1 + \text{SINR}_{u,t,k,r}). \quad (26)$$

In the final post-processing, we consider the effective sum SE (ESSE), which accounts for the overhead due to beam training by removing the corresponding time/frequency resources due to training and feedback of the beam management system $(T_{\text{BM}}, K_{\text{BM}})$ from the spectral efficiency calculation

$$\text{ESSE} = \sum_{u=0}^{U-1} \sum_{\substack{t \notin T_{\text{BM}} \\ k \notin K_{\text{BM}}}} \sum_{r=0}^{R-1} \log_2(1 + \text{SINR}_{u,t,k,r}). \quad (27)$$

With this goal and the beam management framework in mind, we can now define the beamspace processing and machine learning codebook design algorithm.

III. BEAMSPACE-CODEX

In this section, we present the Beamspace-Codex which is a neural network architecture and processing setup that is used to generate the SSB codebooks, with considerations for joint SSB-CSI-RS codebook design left to future work. We start by introducing the beamspace observation that provides a consistent basis for a learning algorithm using dynamic codebooks. We then present the supervised learning formulation that enables convergent and dynamic codebooks. Finally, we define the neural network architecture that enables learning the underlying relationship between the beamspace observation and the




Fig. 2. A timing diagram of the process for codebook learning and evaluation. The AI/ML engine provides new codebooks at each SSB interval for the BS. At each CSI-RS period, the BS transmits pilots using the CSI-RS codebook, which is determined based on the SSB codebook and feedback. The CSI-RS feedback is used to determine the precoder used for data transmission.

SSB codebook. The full methodology—processing, architecture, training, and evaluation—is critical to integrating wireless domain knowledge and physical structure into the learning algorithm.

One of the primary steps in designing an appropriate neural network is formulating the input-output relationship to be conducive and consistent for learning. There are two problems with neural networks within the context of 5G beam management. First, the SSBRI feedback corresponds to a beamformer from a codebook that was used previously, which does not provide a consistent representation or meaning with dynamic codebooks. Second, the number of users is not constant within a cell, so the size of the feedback and data dimension changes over time. Although there are specific ways to overcome dynamic data sizes, it often requires sacrificing hardware optimizations and impedes inference and training times.

We begin by transforming the codebook into the beamspace domain. The beamspace is an angular representation corresponding to the array factor evaluated over a range of directions [34] that has been shown to be beneficial for learning in similar settings [23]. This ensures that regardless of the specific beamforming




Fig. 3. A depiction of the BSC pre-processing, neural architecture, and post-processing. The feedback is converted to an angular representation (beamspace) and the strongest directions are identified based on the feedback reported. The neural network jointly processes the SSB feedback and beamspace representation to produce a new beamspace codebook that is then converted back to complex beamforming coefficients. This format brings together the various codebooks to a consistent input dimensionality even with varying numbers of users or even antenna sizes.

codeword, or even physical antenna dimension, the input from a set of previous beamformers represents a consistent two-dimensional grid of projections. First, we define an angular transformation matrix for antenna size N_s and N_m angular directions as a series of array responses

$$\boldsymbol{\theta}_{N_m} = \frac{1}{\pi} [0, 1, \dots, N_m - 1]^T \quad (28)$$

$$\mathbf{U}_{N_s, N_m} \triangleq [\mathbf{a}_{N_s}(\boldsymbol{\theta}_0), \dots, \mathbf{a}_{N_s}(\boldsymbol{\theta}_{N_m-1})]. \quad (29)$$

The beamspace conversion for N_{x0} azimuth directions and N_{y0} elevation direction is calculated as

$$\mathbf{F}_i^{(N_X \times N_Y)} = [\mathbf{B}_{0:N_X, i}^{\text{SSB}}, \dots, \mathbf{B}_{(N_Y-1):N_X, i}^{\text{SSB}}] \quad (30)$$

$$\mathbf{O}_i^{\text{BSC}} = \mathbf{U}_{N_X, N_{x0}}^* \mathbf{F}_i^{(N_X \times N_Y)} \mathbf{U}_{N_Y, N_{y0}} \quad \forall i. \quad (31)$$

The codebook must be reshaped from a vector of size N_T to the planar dimensions $N_X \times N_Y$ in (30) before the beamspace conversion in (31) that produces the beamspace observation, $\mathbf{O}^{\text{BSC}}$. Note that the beamspace conversion is a reversible operation, so we also train the network to predict the beamspace of the desired output, rather than direct beamforming coefficients. The predicted beamformers are obtained in post-processing, and the computational complexity can be controlled by changing the observation sizes N_{x0} and N_{y0} . In addition to the angular representation, the input is also concatenated with the feedback corresponding to the number of users reporting each beam and the sum RSRP. We can express this using the notation $\mathbf{1}_{(\cdot)}$ to refer to a vector with all 0 entries except the indices where the subscript is true contain a 1 and resulting in the input

$$\mathbf{O}^{\text{BSC}} = \{[\mathbf{O}^{\text{BSC}}[i], \sum \mathbf{1}_{m=i}, \mathbf{1}_{m=i}^T \mathbf{P}]\}_{i=0}^{L_{\max}-1}. \quad (32)$$

With the pre/post-processing defined, we can now ex-

plain the supervised learning environment and desired outputs used for training.

Supervised learning provides many advantages over unsupervised or reinforcement learning when there is a desired output compared to just a black-box metric. The task of determining beamforming vectors that maximize the power to a set of users has an optimal solution if CSI is available via singular value decomposition (SVD) beamforming. To that end, we assume a dataset of channels has been built up or simulated to allow for offline training, although our site-specific results in Section IV suggest that transferring a learned model between sites is still more effective than DFT codebooks. Then, the solution to maximize the power received for a user is the right singular vectors (RSV) of the SVD of the channel [4]. In the MU-MIMO OFDM case, this means calculating the RSVs corresponding to the $L_{\max}$ strongest singular values of all users, assuming perfect CSI is available during the training stage. CSI is not available in a realistic system, however, we train the neural network to learn the mapping from the initial beamspace to the RSV beamspace without CSI.

Additionally, the timing of the beams is important and each beam has a specified set of resources on which to be transmitted. We must sacrifice the time relationship of the data, however, to build a true supervised setting. In particular, we do not use information from the past SSB feedback beyond the previous step. This limitation removes the capability of using historical data as inputs but significantly improves convergence compared to a reinforcement learning setting. The process is restricted to a two-step relationship where we include the previous codebook and feedback to help design the next

codebook. We now have a supervised framework and pre/post-processing step that converts the previous codebook and RSV codebook to the beamspace domain. The final component of the Beamspace-Codex is the neural architecture that facilitates efficient learning.

Determining the best neural network architecture is often challenging even with hyperparameter tuning and an intuitive understanding of the problem. Using the beamspace conversion, the problem can roughly be seen as a translation task from the initial beamspace to the RSV beamspace. It is not necessarily obvious how the limited feedback information could be used to recover the RSV codebook. There is some underlying information available, however, by the selection of one beam over $L_{\max} - 1$ other options. Additionally, the environment and user densities follow a pattern that is also possible to exploit. At the same time, the inputs and outputs are similar to an image of the angular directions of transmitted power. Therefore, we considered traditional architectures (fully-connected networks), image-based architectures (convolutional and diffusion [35] networks), and translation-based architectures (transformers [36] and autoencoders). Although the problem has many similarities with both vision and translation problems, state-of-the-art architectures were not necessarily best under reasonable computational limits i.e. 24GB of VRAM and 5ms inference time. Throughout our testing, the fully connected architecture was typically better performing, however, with sufficient hyperparameter tuning we found convolutional networks had the potential to capture and learn the beamforming codebook more effectively while being computationally efficient.

We propose a multi-layer convolutional network with pooling and dropout comprised of with parameters and sizes determined through Bayesian hyperparameter tuning and defined in Table I. A comparison of model architectures is shown in Figure 7 within Section V. The model is trained using an Adam optimizer [37] with early stopping and cosine learning rate decay. Training is performed with a 550,000 sample training dataset following the setup in Section V with validation and test sets each from an unseen portion of the data corresponding to 160,000 and 80,000 samples respectively. Gradients are determined by the partial derivative of the cosine distance ($d_{\cos}(\mathbf{a}, \mathbf{b}) = 1 - \mathbf{a}^* \mathbf{b} / (\|\mathbf{a}\| \|\mathbf{b}\|)$ [25]) between the RSV beamspace and the neural network-predicted beamspace for each neural network weight. The choice of the loss function is critical for performance. For example, the same network architectures trained on mean squared error results in significantly worse ($> 5\text{dB}$) performance due to focusing on the pixel-wise magnitude. In comparison, the cosine distance focuses on the relative magnitude or direction of the RSV beamspace.

IV. DATA GENERATION SETUP

We require an accurate channel simulator coupled with a flexible framework that retains the timing of the 5G beam training process to fairly evaluate the performance of ML-designed codebooks. We integrate the channel generation from QuaDRiGa [26] with a post-processing suite that enables an initial access scenario, beamspace generation, codebook determination, and evaluation of the network performance after accounting for beam management overhead. QuaDRiGa generates channels through a stochastic process with additional features for spatial consistency, correlations, and spherical wave modeling. In our simulations, we simulate 200 users scattered over a BS sector site. The BS is equipped with an $N_X = 4$, $N_Y = 8$ planar array using a 3GPP 3D antenna model with half-wavelength spacing. Each user has $N_R = 4$ antennas all tuned for a 3.5GHz carrier frequency as might be found in modern UEs which are typically capable of 4 layer MIMO [38]. 10% of users are given a vehicular mobility pattern such that they travel along a roadway through the center of the cell with speeds normally distributed with mean 25m/s and variance 5m/s. The remaining users are uniformly scattered within 450m and travel in any direction with speeds uniformly drawn from an interval of $[0, 3]\text{m/s}$. Channel distributions follow a 3GPP urban macrocell environment and channels are sampled over a bandwidth of 100MHz every 1ms for 2s. A large number of users are simulated to build up a database that is spatially and environmentally consistent. In the final section of the results, we address how the model generalizes to a new setting with different environmental parameters and different roadway locations.

Once the channels are generated, post-processing is necessary to produce a realistic setting and build up the algorithm's datasets. The initial access procedure starts by randomly selecting a number of active users (uniformly between $[4, 12]$) and gathering the corresponding channels and time slots from the channel set. Then uniformly randomly selecting $L_{\max}$ DFT beams for the 'prior' codebook and calculating the RSV codebook. The codebooks are converted to beamspace representations and stored in a dataset along with the SSB feedback from a subset of the active users where each active user is assigned an 80% probability of its feedback being included. This represents an initial access channel where not all users are previously known and introduces an additional regularization term into our algorithm to prevent over-focusing on known users. It is assumed that UEs use antenna selection for the RSRP reception during the SSB process, digital combining for CSI-RS, and LMMSE combining during data transmission as defined in the system model.

TABLE I
BEAMSPACE-CODEX NETWORK PARAMETERS

Layer	Primary Parameter	Activation	Output Dimension
Conv+Max Pool	128 Filters	ReLU	$\lceil (N_{x0} + 2)/2 \rceil \times \lceil (N_{y0} + 2)/2 \rceil \times 128$
Conv+Max Pool	96 Filters	ReLU	$\lceil (N_{x0} + 2)/4 \rceil \times \lceil (N_{y0} + 2)/4 \rceil \times 96$
Dropout	0.3 Rate		$\lceil (N_{x0} + 2)/4 \rceil \times \lceil (N_{y0} + 2)/4 \rceil \times 96$
Conv	320 Filters	ReLU	$\lceil (N_{x0} + 2)/4 \rceil \times \lceil (N_{y0} + 2)/4 \rceil \times 320$
Flatten			$(\lceil (N_{x0} + 2)/4 \rceil)(\lceil (N_{y0} + 2)/4 \rceil)(320)$
Dropout	0.5 Rate		$(\lceil (N_{x0} + 2)/4 \rceil)(\lceil (N_{y0} + 2)/4 \rceil)(320)$
Fully Conn.	$2N_{x0}N_{y0}L_{\max}$ Neurons		$2N_{x0}N_{y0}L_{\max}$
Reshape			$N_{x0} \times N_{y0} \times 2L_{\max}$

V. SIMULATION RESULTS

It is critical to extensively evaluate a neural network model to fairly represent and compare it against traditional methods. In our previous work [39], we have considered a learning threshold as a minimal condition for determining if meaningful relationships have been learned by the network. In realistic wireless channels, such settings are much harder to characterize, so we instead benchmark the performance against industry-standard DFT techniques and RSV codebooks assuming CSI were known perfectly. RSV codebooks identify the upper-performance limits (with respect to RSRP-maximizing) and DFT codebooks identify the minimal performance needed to justify moving beyond traditional codebooks. We then use our simulation framework to better understand how the various codebooks and type-II CSI feedback parameters affect network performance.

A. RSRP and SNR

The RSRP is a basic metric that can be used to directly demonstrate the ability of the neural network to output a valuable SSB codebook. Following the SSB process, we also highlight the performance that the dynamic codebook provides when selecting CSI-RS beams as a decomposition of the SSB codebook. While these two metrics—RSRP and SNR—do not ultimately characterize the network performance as a result of the codebook algorithm, they do provide meaningful insight into the performance of various codebook methods. Figure 4(a) shows the empirical cumulative distribution function (CDF) of the RSRP reported using various codebook/beamforming methods compared to our proposed BSC algorithm. Without any beamforming, more than 50% of users would be below the minimum RSRP for signal detection, which is often around -120dBm . It can also be seen that the proposed codebook does not negatively impact worst-case users, with about 10% of users falling under the sensitivity threshold compared to

over 20% of users with DFT codebooks. To showcase the extensibility of the proposed BSC algorithm, we also include results using $N_T = 16$ antennas without changing the model, preprocessing, or retraining. The performance is unchanged and actually improves versus the ideal RSV results, even though the network was trained with data from a larger number of antennas. This presents a significant advantage for network operators that can immediately deploy different antenna arrays without retraining or gathering new data.

The SSB codebooks are decomposed into CSI-RS codebooks in Figure 4(b), where we see the RSV codebook is *not* consistently the best SSB codebook when directly decomposed into CSI-RS codebooks, although the decomposed BSC is not universally better than decomposed DFT codebooks either. We can understand this as a result of the rich scattering environment that causes the RSV to not effectively decompose into a small number of incoherent DFT components. While RSV beamformers maximize the received power for a user with perfect CSI, this maximization is based on the signal coherently combining and is not well-represented by a narrow DFT beam. From these results, we can see that the BSC algorithm is advantageous compared to traditional methods, but we also see that restricting the CSI-RS codebook to oversampled DFT beams is a significant limitation.

B. CSI-RS reporting

In this subsection, we provide an in-depth evaluation of the CSI-RS process without BSC codebooks. We evaluate performance based on the effective downlink SE after overhead due to beam training, uplink feedback, and an additional signaling overhead factor of 10%. Uplink feedback is assumed to be the only resource allocated for uplink transmission, and we assume a worst-case scenario where users transmit with MCS0 from Table 5.1.3.1-1 [27]. User selection during the data phase is performed based only on the knowledge at the BS by




Fig. 4. Two plots of the codebook performance measuring (a) SSB-RSRP during 5000 random active user groups and (b) the sliding window average SNR during CSI-RS reception. During the SSB stage, each of the codebooks is used directly, and then in (b) the SSB codebooks are decomposed to form the active CSI-RS codebook subset. Non-beamformed transmission presents a significant loss in performance shown by the “gain of BF” arrow, while the learned BSC codebook shows gains of over 3dB compared to traditional DFT beams. BSC also seamlessly applies to smaller N_T with the dotted lines showing the relative performance vs RSV actually improving with less antennas than trained with. While RSV beamforming is an upper bound on the RSRP, the DFT decomposition of an RSV codebook is not necessarily SNR-maximizing due to multipath propagation.

first estimating the sum spectral efficiency achieved with the imperfect CSI for all users. The best estimated SE combination is selected and used until the next SSB cycle. Such a format is naturally suboptimal, even with perfect CSI, due to a lack of rank and power adaption. User selection is not the focus of this investigation, so we only use an exhaustive search over the reconstructed channel knowledge available at the BS to provide a reasonable scheduling algorithm. The SE is also compared with rank-1 non-PMI beamforming where RZF is not employed, the channel is not estimated, and the only feedback is obtained from the CRI, with user selection and precoder determination simply being the CSI-RS beamforming codewords with the strongest reported user RSRP. This is often an envisioned strategy for sparse mmWave channels with low-latency feedback but is not assumed to be effective in sub-6GHz channels due to the rich scattering environment.

The size of the CSI-RS codebook is related to the channel estimation SNR because larger CSI-RS codebooks, made of orthogonal DFT beamformers, will result in equivalent or higher SNR compared to smaller codebooks with the same channels and feedback. We characterize the performance of larger or smaller CSI-RS codebooks in Figure 5(a). Although the estimation SNR will always improve with a larger codebook, there is also a modest amount of overhead with larger P_{CSI} so that there is little spectral efficiency gain beyond $P_{\text{CSI}} \geq 8$.

In the next comparison, the feedback quantization is evaluated by modifying the parameter L_{CSI} in type-II formats. We can see from Fig. 5(b) that higher resolution feedback is beneficial, with the effective spectral efficiency exceeding 2 times higher when the resolution is improved from $L_{\text{CSI}} = 1$ to $L_{\text{CSI}} = 32$. From there on, the other parameters—CSI-RS resource block allocation (NRB) and the number of frequency selective resources (S_B)—are less influential compared to the quantization resolution. The modest performance gains achieved with multiple frequency selective precoders and larger downlink CSI-RS resource usage are offset by the increasing overhead resulting in negligible or even performance loss. Joining the information from Figures 5(a)-6(a), a good balance of feedback and performance can be achieved with a setting of $L_{\text{CSI}} = 32$, $S_B = 1$, $\text{NRB} = 24$, $P_{\text{CSI}} = 8$. It is important to note that the current 5G Release 16, even with enhanced type-II feedback, only supports $L_{\text{CSI}} \leq 6$, which means most situations will only yield about 1.5 times higher effective spectral efficiency relative to SU-MIMO and nearly equivalent performance to the non-PMI beamforming using only SSB codebooks. Perhaps even more surprising is that using a single CSI-RS beam ($P_{\text{CSI}} = 1$) has less degradation than using $L_{\text{CSI}} = 8$ feedback beams per rank, relative to the best performance. This suggests that feedback limitations are more likely to restrict MU-MIMO performance in sub-6GHz systems compared to codebook optimization or accurate channel estimation.




Fig. 5. Evaluation of the effects of active CSI-RS codebook size (a) and quantization using L_{CSI} feedback components (b). We can use the size of the codebook as a means of looking at how codebook optimization improves performance because the overhead with larger P_{CSI} is very small compared to the overhead of type-II feedback. SU-MIMO results using type-I PMI are shown by dashed lines while MU-MIMO results are shown by solid lines and Non-PMI refers to rank 1 beamforming using the CRI. SU-MIMO does not include multipath feedback so only a single SU-MIMO result is shown in (b). Other parameters are set to $L_{\text{CSI}} = 32$ for (a), $P_{\text{CSI}} = 16$ for (b), $S_B = 1$, $\text{NRB} = 24$.




Fig. 6. A comparison of the effective spectral efficiency with different numbers of frequency selective precoders S_B and varying the number of resource blocks (NRB) per S_B . It is assumed that $L_{\text{CSI}} = 32$, $P_{\text{CSI}} = 16$, $\text{NRB} = 24 \times S_B$ in (a) and $S_B = 1$ in (b). While increasing the number of resource blocks, therefore increasing the number of pilots, also increases the channel estimation performance, the rising feedback overhead causes the performance to decrease in (b). In comparison, the number of S_B increases the overhead linearly, yet the performance is essentially equivalent in (a). Therefore we see that frequency selective feedback and beamforming is not necessarily advantageous when overhead is accounted for.

C. Neural architectures

Looking at the BSC neural network, it remains to be determined whether the proposed architecture is a good choice of neural network for the problem. Although the beamspace inputs are more closely aligned with an image-processing task, where convolutional neural networks (CNN) are advantageous, we found that performance varied significantly and hyperparameter tuning

was critical. We believe this is due to the lack of shift-equivariance which is a core learning bias of CNNs. For example, if a region of the beamspace that is “active” is shifted, does not necessarily mean that the RSV codebook will shift in a similar way. This is especially true because the beamspace does not include newly-active UEs, similar to an initial access scenario. Therefore appropriate regularization via parameter tuning is nec-




Fig. 7. A comparison of the RSRP performance of classic neural architectures—multilayer perceptron (MLP), convolutional autoencoder (AE)—within the beamspace framework. It can be seen that although the architectures have comparable results, the BSC architecture based on CNNs is the best choice for learning beamspace features.

essary. Figure 7 shows the proposed BSC network outperforming different architectures trained with the same data, optimization strategy, and comparable execution time by about 1dB on average.

D. Site generalization

Up to now, we have used a test set of data drawn from the same statistical environment as the training set. While a network operator will likely gather data for a specific site, here we examine the impact of testing on a new distribution of data. We simulate a new environment with completely different user distributions and mobility patterns. Now 50% of users are vehicular with roadways defined in new locations and a new stochastic realization of the channel environment. In Figure 8 the agnostic (orange) data is obtained without updating the model, whereas the fine-tuned data (blue) shows the results after we allow for 60s of retraining on data from the new environment. Our time limit is arbitrary and could ultimately be orders of magnitude shorter than the timescale a network operator might use to update the model, but corresponds to approximately 1% of the training time used in our setup for a balance of fine-tuning and generalization. Figure 8 shows the performance in the new environment (agnostic) is noticeably worse relative to the RSV. The performance loss is improved during fine-tuning, especially for the lowest 10% of users.

VI. CONCLUSION

In this paper, we presented a novel framework for learning dynamic codebooks for sub-6GHz 5G NR. We set up




Fig. 8. Histogram representing the performance delta between BSC and RSV codebooks in the agnostic setting and after fine-tuning. It can be seen that the worst-case differences are resolved and in some cases (when more users are present than SSB beams) the RSV codebook can be outperformed based on the higher proportion of the blue columns to the right of the orange columns.

a system with multiple UEs each equipped with a uniform linear array to receive from a BS with a fully digital planar array. We simulate realistic sub-6GHz channels using QuaDRiGa and build a post-processing framework for evaluating the system performance of SSB and CSI-RS codebooks and feedback. Using this framework we address two questions related to codebook performance in 5G: 1) Can machine learning design more effective SSB codebooks? 2) How should feedback be configured in sub-6GHz environments?

We first proposed a codebook transformation based on beamspace projections that allowed for efficiently and consistently representing SSB codebooks in a format conducive to learning. This beamspace representation can circumvent challenges with dynamic user numbers, changing codebooks, and even deploying across different array sizes. Using the beamspace as a translation tool, we designed a neural architecture, Beamspace-Codex, to output new codebooks from SSB feedback. The proposed method directly integrates with current 5G standards and improves SSB RSRP by 2-3dB on average. We also show that only 1% of the training compute budget is needed for retraining in new environments to significantly improve the worst-case performance. Secondly, we found that the PMI quantization resolution, determined by the number of feedback DFT beams, is a limiting factor

in the performance of MU-MIMO in current standards. In future work, we will introduce machine learning throughout the beam management process.

REFERENCES

- [1] R. M. Dreifuerst and R. W. Heath, "Massive MIMO in 5G: How beamforming, codebooks, and feedback enable larger arrays," in *submission to IEEE Commun. Mag.*, 2023. [Online]. Available: <https://arxiv.org/abs/2301.13390>
- [2] K. Lee *et al.*, "Extreme massive MIMO for upper-mid band 6G communications," in *International Conference on Information and Communication Technology Convergence*, 2022, pp. 997–999.
- [3] G. Morozov, A. Davydov, and V. Sergeev, "Enhanced CSI feedback for FD-MIMO with beamformed CSI-RS in LTE-A Pro systems," in *Proc. of the IEEE Veh. Technol. Conf.*, Jul. 2016.
- [4] V. Raghavan *et al.*, "Beamforming Tradeoffs for Initial UE Discovery in Millimeter-Wave MIMO Systems," *IEEE J. Sel. Top. Signal Process.*, vol. 10, no. 3, pp. 543–559, Apr. 2016.
- [5] Y. Heng *et al.*, "Six key challenges for beam management in 5.5G and 6G systems," *IEEE Commun. Mag.*, vol. 59, no. 7, pp. 74–79, Jul. 2021.
- [6] M. Giordani *et al.*, "A tutorial on beam management for 3GPP NR at mmwave frequencies," *IEEE Commun. Surv. Tutorials*, vol. 21, no. 1, pp. 173–196, Jan. 2019.
- [7] A. Alkhateeb *et al.*, "Channel estimation and hybrid precoding for millimeter wave cellular systems," *IEEE J. Sel. Top. Signal Process.*, vol. 8, no. 5, pp. 831–846, Oct. 2014.
- [8] E. Becirovic, E. Bjornson, and E. G. Larsson, "Combining reciprocity and CSI feedback in MIMO systems," *IEEE Trans. Wirel. Commun.*, May 2022.
- [9] Samsung, "Massive MIMO for new radio," Dec. 2020. [Online]. Available: <https://www.samsung.com/global/business/networks/insights/white-papers/1208-massive-mimo-for-new-radio/>
- [10] B. Li *et al.*, "On the efficient beam-forming training for 60GHz wireless personal area networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 2, pp. 504–515, 2013.
- [11] V. Va *et al.*, "Inverse multipath fingerprinting for millimeter wave V2I beam alignment," *IEEE Trans. Veh. Technol.*, vol. 67, no. 5, pp. 4042–4058, May 2018.
- [12] Y. Wang *et al.*, "Site-specific online compressive beam codebook learning in mmWave vehicular communication," *IEEE Trans. Wirel. Commun.*, pp. 1–14, 2021.
- [13] Q. Xue *et al.*, "Beam management in ultra-dense mmWave network via federated reinforcement learning: An intelligent and secure approach," *IEEE Trans. Cogn. Commun. Netw.*, 2022.
- [14] J. Yang, W. Zhu, and M. Tao, "Deep learning for hierarchical beam alignment in mmWave communication systems," in *Proc. of the IEEE Global Communications Conference*, 2022.
- [15] J. Wang *et al.*, "Beam codebook based beamforming protocol for multi-gbps millimeter-wave WPAN systems," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1390–1399, 2009.
- [16] Z. Xiao *et al.*, "Hierarchical codebook design for beamforming training in millimeter-wave communication," *IEEE Trans. Wireless Commun.*, vol. 15, no. 5, pp. 3380–3392, 2016.
- [17] R. Shafin *et al.*, "Self-tuning sectorization: Deep reinforcement learning meets broadcast beam optimization," *IEEE Trans. Wireless Commun.*, vol. 19, no. 6, pp. 4038–4053, Jun. 2020.
- [18] W. Xia *et al.*, "A deep learning framework for optimization of MISO downlink beamforming," *IEEE Trans. Commun.*, vol. 68, no. 3, pp. 1866–1880, Mar. 2020.
- [19] Y. Heng, J. Mo, and J. G. Andrews, "Learning site-specific probing beams for fast mmwave beam alignment," *IEEE Trans. Wirel. Commun.*, vol. 21, no. 8, pp. 5785–5800, 2022.
- [20] R. M. Dreifuerst, R. W. Heath, and A. Yazdan, "Massive MIMO beam management in sub-6 GHz 5G NR," in *Proc. of the IEEE Veh. Technol. Conf.*, 2022, pp. 1–5.
- [21] L. Li *et al.*, "CSI sensing from heterogeneous user feedbacks: A constrained phase retrieval approach," *IEEE Trans. Wireless Commun.*, Feb. 2023.
- [22] K. Li *et al.*, "Learning enhanced beamforming vector from CQIs in 5G NR FDD massive MIMO systems: A tuning-free approach," in *Proc. of the IEEE Workshop on Signal Processing Advances in Wireless Communications, SPAWC*, Sep. 2021.
- [23] C. K. Wen, W. T. Shih, and S. Jin, "Deep learning for massive MIMO CSI feedback," *IEEE Wirel. Commun. Lett.*, vol. 7, no. 5, pp. 748–751, 2018.
- [24] J. Kim, H. Lee, and S.-H. Park, "Learning robust beamforming for MISO downlink systems," *IEEE Communications Letters*, vol. 25, no. 6, pp. 1916–1920, 2021.
- [25] M. Chen *et al.*, "Deep learning-based implicit CSI feedback in massive MIMO," *IEEE Trans. Commun.*, vol. 70, pp. 935–950, Dec. 2022.
- [26] S. Jaeckel *et al.*, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, pp. 3242–3256, 2014.
- [27] 3GPP, "TS 38.214 V16.7.1 NR; Physical layer procedures for data."
- [28] A. Singh and S. Joshi, "A survey on hybrid beamforming in mmWave massive MIMO system," *J. Sci. Res.*, vol. 65, no. 01, pp. 201–213, 2021.
- [29] R. W. Heath Jr. and A. Lozano, *Foundations of MIMO Communication*. Cambridge University Press, 2018.
- [30] L. Godara, "Application of antenna arrays to mobile communications. ii. beam-forming and direction-of-arrival considerations," *Proceedings of the IEEE*, vol. 85, no. 8, pp. 1195–1245, 1997.
- [31] M. R. Castellanos *et al.*, "Channel-reconstruction-based hybrid precoding for millimeter-wave multi-user MIMO systems," *IEEE J. Sel. Areas Commun.*, vol. 12, no. 2, pp. 383–398, May 2018.
- [32] MathWorks, *5G Toolbox*, 2023. [Online]. Available: <https://www.mathworks.com/help/5g/>
- [33] R. W. Heath *et al.*, "An overview of signal processing techniques for millimeter wave MIMO systems," *IEEE J. Sel. Top. Signal Process.*, vol. 10, no. 3, pp. 436–453, Apr. 2016.
- [34] A. Sayeed, "Deconstructing multiantenna fading channels," *IEEE Trans. Signal Process.*, vol. 50, no. 10, pp. 2563–2579, 2002.
- [35] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," 2020. [Online]. Available: <https://arxiv.org/abs/2006.11239>
- [36] Z. Liu *et al.*, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. of the IEEE Int. Conf. Comput. Vis.*, Mar. 2021, pp. 9992–10002.
- [37] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, Dec. 2014.
- [38] J. Yang, X. Ge, and Y. Zhong, "How much of wireless rates can smartphones support in 5G networks?" *IEEE Network*, vol. 33, no. 3, pp. 122–129, 2019.
- [39] R. M. Dreifuerst and R. W. Heath, "SignalNet: A low resolution sinusoid decomposition and estimation network," *IEEE Trans. Signal Process.*, vol. 70, pp. 4454–4467, Aug. 2022.