ABSTRACT

DREIFUERST, RYAN MITCHELL. Machine Learning-Assisted Beam Management and Codebook Design. (Under the direction of Robert W. Heath jr.).

Antenna arrays capable of beamforming and spatial multiplexing have played a foundational role in standardized wireless systems. Scaling up these arrays enables greater degrees of flexibility and resource optimization, but maximizing the performance requires accurate channel station information (CSI). Obtaining CSI for large multiple-input multiple-output (MIMO) arrays requires overhead in the form of pilot transmission, beam training, and feedback. Modern mobile broadband supports this process through a strategy called *beam management* that combines two tiers of codebook-based beam training with quantized feedback. Yet these beamforming codebooks introduce a source of inefficiency, especially if they do not capture environment characteristics.

In this dissertation, we propose a series of algorithms that integrate machine learning into the codebook design process. First, we introduce the concept of beamspace learning, which forms a core element of the codebook learning paradigm. The beamspace is used to dynamically configure the initial access beam training codebook without modifying the other codebooks. In this way, the proposed algorithm is directly supported by current 5G specifications and can be used to improve beam training performance. In the second step, we expand the formulation to jointly configure both beam training codebooks to support extremely large arrays for 6G. In this setting, we design an end-to-end framework that optimizes the codebooks together to achieve the highest beam alignment. Our results show nearly optimal beam training performance with extremely large arrays. Finally, we propose an end-to-end learning framework for interference-aware, rate-maximizing codebook design across a network. In each of these works, we focus on the system-level impact

to characterize how the novel designs translate to spectral efficiency improvements. Our results show valuable improvements in terms of beam alignment, outage probability, and effective sum spectral efficiency across a range of scenarios and geometries. Based on this work, machine learning can effectively improve codebook design in beam management.

Machine Learning-Assisted Beam Management and Codebook Design

by
Ryan Mitchell Dreifuerst

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Electrical Engineering

Raleigh, North Carolina
2024

APPROVED BY:

Nuria González-Prelcic	Alexandra Duel-Hallen
Yuchen Liu	Robert W. Heath jr. Chair of Advisory Committee

DEDICATION

This work is dedicated to my wife, Maddie, who has supported me as we moved across the country twice to pursue this dream. I am eternally grateful for her help proofreading papers, supporting me during deadlines, and her endless patience in hearing about the latest ideas occupying my brain. I also dedicate this work to my family who encouraged my love for discovery from a young age. Together they have been the perfect support and encouragement

BIOGRAPHY

Ryan Dreifuerst graduated with a B.S. in Electrical Engineering from Milwaukee School of Engineering (MSOE) and a B.S. in Electrical and Communications Engineering from Technische Hochschule Lübeck. He received a M.S. from the University of Texas at Austin where he researched low-resolution multi-sinusoidal decomposition and estimation methods. Ryan's research interests lie at the intersection of wireless communications, signal processing, and machine learning. In his undergraduate, Ryan was the top-ranked student at MSOE and was the recipient of three awards for academic and community excellence. In graduate school, his work on ML for wireless was awarded a multi-year NSF grant. Ryan has spent the previous summers with Meta, Qualcomm, Samsung Research America, and Nvidia researching and advancing novel algorithms for wireless systems. Currently, his work is focused on deep learning for massive MIMO systems, including codebook design for beam management and feedback.

ACKNOWLEDGEMENTS

First, I would like to thank my family for their unending support. My wife, Maddie, has walked every step of this journey with me and truly has been a constant source of support. She has been steadfast through cross-country moves to Texas and North Carolina, supportive during all the late nights working on papers, and patient through every technical challenge and joy I have gone through over these years. In many ways, we would joke that she has been a graduate student right alongside me for all the ways she has also had to sacrifice and overcome challenges with me. It is clear to me that being married to a graduate student is like being a graduate student just without your name on the diploma. I also want to thank my parents and brothers who have been another source of support throughout my education, providing encouragement and consolation whenever needed. These last five years of weddings, holidays, and many long drives have been a testament to the ways my family has supported me in my PhD.

Second, I would like to thank my supervisor, Prof. Robert Heath, for his guidance over the last 5 years. He has given me the tools and the skills to bring ideas to fruition and work through challenges. He has helped me to become a better researcher, writer, and presenter while providing opportunities for me to engage and learn from professionals throughout our field. I would like to also thank my committee members: Prof. Nuria González-Prelcic, Prof. Alexandra Duel-Hallen, and Prof. Yuchen Liu for their help and technical discussions. Also, additional thanks are due to Prof. Aritra Mitra who helped serve as a substitute committee member during my defense.

Third, I would like to thank all of my friends, lab mates, and coworkers who have helped me get here. My friends in Austin deserve extra gratitude for spending so many hours in front of a whiteboard with me. I will fondly remember the days spent grinding away on homework and commiserating in the Garage together. To my many friends in Raleigh who have given me advice, muddled through ideas together, or simply been a welcoming community–thank you for being exceptional people to do life with.

Finally, I would like to thank God who has been gracious, consistent, and good at every moment. My faith is integral to how I see the world, and I am grateful that I have been blessed with so many opportunities. It always struck me to see His creation as an act of engineering. Pursuing Jesus and His plan is the cornerstone that my life is built upon.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

INTRODUCTION

In this chapter, we provide context and background for the dissertation. We motivate the use of large-scale arrays in mobile broadband to enable beamforming and spatial multiplexing. We then introduce beam management as it is specified in 5G New Radio (NR) as an enabling framework for massive MIMO. Then, we discuss the role of machine learning for physical layer signal processing in wireless systems. In the final section, we provide an overview of the contributions of this dissertation and outline the organization of the remaining chapters.

1.1 Massive arrays in cellular systems

Large antenna arrays are envisioned as a cornerstone to support new bands and higher data rates in next-generation mobile broadband (Morozov et al. 2016). Historically, cellular systems operated at low carrier frequencies, over relatively small bandwidths, and employed low dimensional multiple-input-multiple-output (MIMO) (Dreifuerst and Heath 2023b). As a result, multi-antenna operations were typically employed for spatial multiplexing rather than for beamforming or null shaping (Li et al. 2016) (Godara 1997). At higher carrier frequencies, however, the effective aperture of an antenna shrinks and larger bandwidths become available, reducing the wireless link budget (Heath et al. 2016). To improve the wireless link, multiple antennas can be configured together to coherently transmit and receive the signal. Effectively employing antenna arrays, however, requires accurate channel state information.

Cellular systems have incorporated a codebook-based approach to CSI acquisition and array configuration (Giordani et al. 2019; Dreifuerst and Heath 2023b). Codebook-based methods reduce the amount of feedback required for obtaining CSI by representing the feedback as indices from a set of options defined as the "codebook" (Qin and Yin 2023). The codebooks can be blind, referring to a set of options that are not known by both the transmitter and receiver, or shared when used to encode and decode information. In practice, modern cellular systems implement CSI acquisition through multiple blind codebooks for beam training and a shared codebook for CSI quantization. This unified process is termed *beam management* and serves a fundamental role in the integration of Massive MIMO arrays in cellular systems.

1.2 Beam management

Beam management as specified in 3GPP Releases 15-18 can be understood by five distinct stages: initial access, beam refinement, CSI feedback, data transmission, and beam tracking (Dreifuerst and Heath 2023b). From an array-processing perspective, these stages each serve different roles in configuring and updating the wireless link. Initial access serves as a preliminary stage where users can join the network and determine a coarse beamforming vector. Beam refinement attempts to improve the link by iterating over a subset of precise

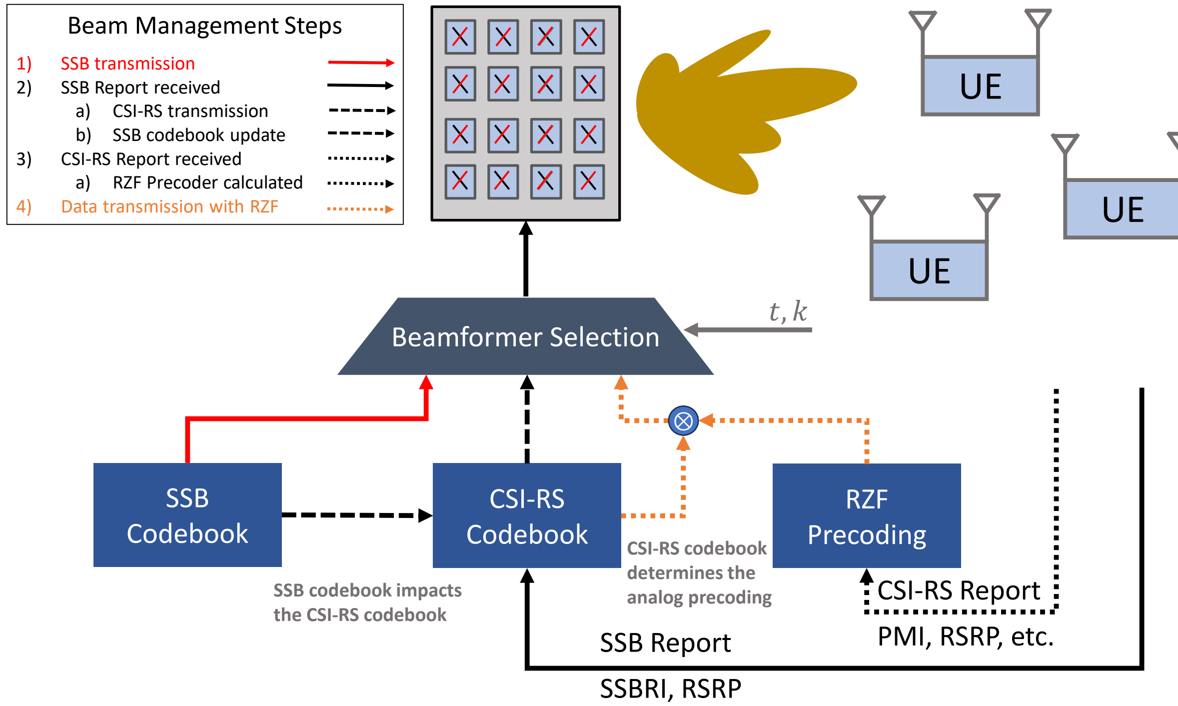

Figure 1.1: A depiction of the interaction between the codebooks and data transmission in beam management. The SSB and CSI-RS codebooks are used to determine an effective analog precoding that is combined with the regularized zero-forcing (RZF) precoding determined by the CSI-RS report.

beamformers based on the users' initial access feedback (Giordani et al. 2019). The users then report a packet of feedback containing CSI and beamformer selection. Based on these three stages, data transmission involves combining the available CSI to serve one or more users. Finally, beam tracking is used as the users move to reduce overhead when updating beam selection (Va et al. 2016). Beam tracking is not a focus of this work, but the preceding 4 stages are shown in Figures 1.1 and 1.2.

This unified process of beamforming, quantization, and feedback results in interrelated information making the codebook optimization challenging. For example, modifying the first beamforming codebook influences how the second set of beamformers should be chosen, which influences the channel estimation, which impacts the CSI accuracy and

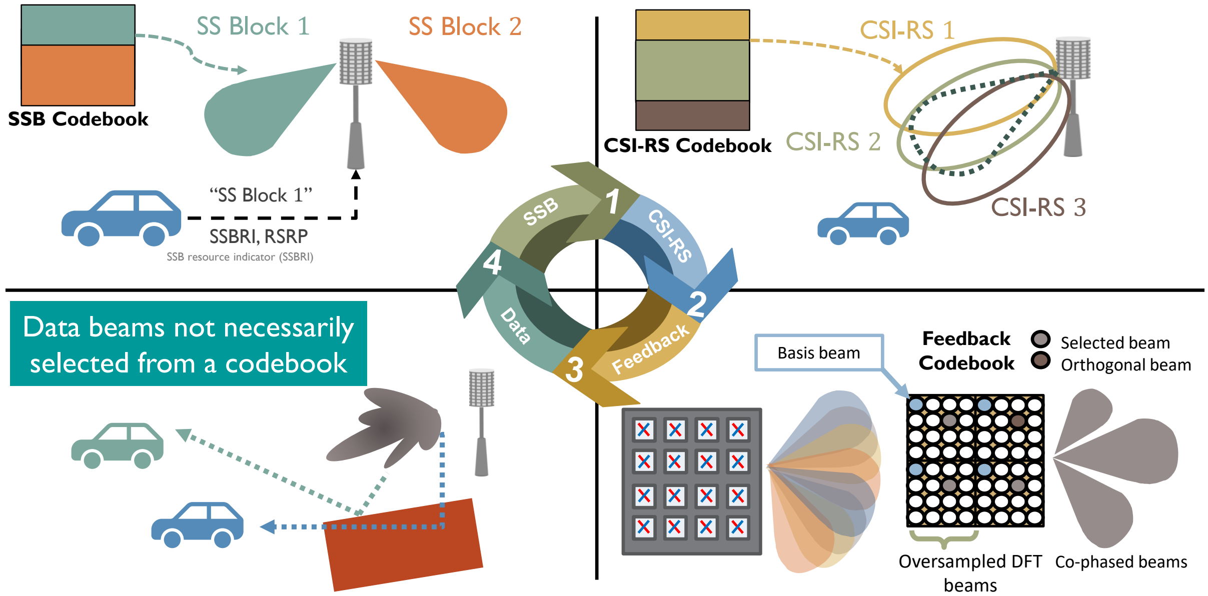

Figure 1.2: A graphical representation of the beam management process timing. The process is iterative and highlights the different codebooks involved at each step which culminate in the data transmission as a combination of the information gathered during all beam management stages.

therefore the achievable data rate. Any algorithm that determines how beam training should be updated is ultimately dependent on the codebooks employed. This multi-step process makes codebook design and optimization require a system-level perspective when optimizing.

1.3 AI in MIMO

MIMO link configuration–from beam training to channel estimation–is an exciting paradigm for artificial intelligence/machine learning (AI/ML) due to the large dimensionality and difficulty of mathematically analyzing (Jiang et al. 2017; Goutay et al. 2020). In fact, significant work has focused on these issues (Sim et al. 2018; Wang et al. 2009; Choi et al.

2016; Xue et al. 2023) and references therein. At the same time, deploying learning algorithms in wireless systems is challenging due to the timing constraints, power/compute budgets, and lack of coherence between transmitter and receiver (Chen et al. 2023). Ultimately, AI/ML has been identified as a potential innovation for 6G, but not for traditional physical layer tasks. Instead, the key areas of focus target beam management, CSI acquisition, and integrated sensing and communications (ISAC) (Lin 2023).

In contrast to classical physical layer tasks like channel estimation, most of the beam management operations have significantly longer periods that can enable deep learning or data-driven algorithms (Dreifuerst and Heath 2023b; Xue et al. 2023; Wang et al. 2017; Zhang et al. 2019). In addition, beam management poses a significant challenge for explicit information modeling through the selection (and non-selection) of a given beamformer. Deep learning is well-suited to capturing this implicit information along with environmental and user density distributions through gradient-based pattern matching. By further incorporating domain knowledge in the form of array processing, beam management is a system that is naturally inclined to be optimized through deep learning.

1.4 Overview of contributions

Beam management is a relatively new process that is characterized by a fundamental tradeoff between overhead, CSI accuracy, and performance. The performance of beam-based systems is dependent on the codebooks involved, which are a source of inefficiency with generic or poorly optimized codebooks. This research addresses the challenge of codebook design through the lens of AI/ML. The contributions of this dissertation are summarized as follows.

- In Chapter 2 we introduce the multi-tier beam management process for a sub-

6GHz 5G system. Within this system, we characterize the roles of each codebook, propose a fast feedback selection algorithm, and introduce the concept of beamspace learning. We use this formulation to design *Beamspace-Codex* a preprocessing and neural architecture that generates dynamic SSB codebooks. The SSB performance is notably improved by the site-specific modeling and capture of multipath scattering in the environment.

- In Chapter 3 We propose an integration of beam management and traditional hybrid array operations for 6G standardization. In this formulation, jointly-learned SSB and CSI-RS codebooks are used to configure the analog portion of hybrid arrays while user-provided feedback is used to configure the digital precoding. We build upon the beamspace learning concept to design *X-BM*, a joint codebook design and end-to-end learning framework for optimizing beam training. This extension improves CSI-RS reception and channel estimation by more than 6dB on average.

- Finally, in Chapter 4 we investigate learned codebooks for inteference-aware multi-cell systems. With multiple base stations serving overlapping coverage regions, the multi-cell system incorporates the effects of interference on beam management. Our proposal elevates end-to-end learning to encompass the achievable spectral efficiency. As a result, the learned codebooks are trained to maximize the achievable rate through the hierarchical beam search and multi-cell interference. The resulting *NBL* algorithm is capable of improving network spectral efficiency by over 25% in dense 6G systems.

1.5 Notation

We use the following mathematical notation throughout this thesis: $\mathbf{A}$ is a matrix, $\mathbf{a}$ and $\{a[i]\}_{i=1}^{N}$ are column vectors and a, A denote scalars. $\mathbf{A}^T$, $\overline{\mathbf{A}}$, $\mathbf{A}^*$, and $\mathbf{A}^\dagger$ represent the transpose, conjugate, conjugate transpose, and psuedo-inverse of $\mathbf{A}$. The real and imaginary parts of $\mathbf{A}$ are denoted by $\mathbb{R}(\mathbf{A})$ and $\mathbb{I}(\mathbf{A})$. $\mathbf{A}[k,\ell]$ 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 $\|\mathbf{a}\|_2 = \mathbf{a}^*\mathbf{a}$ for vectors and the Frobenius norm $\|\mathbf{A}\|_F = \sqrt{\text{Tr}(\mathbf{A}\mathbf{A}^*)}$ for matrices. We use the notation $\lfloor\cdot\rceil$ for rounding to the nearest integer. We define $\mathsf{j} = \sqrt{-1}$. 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 c and u to refer to a specific base station cell and UE, t as a specific time, k as a specific frequency resource, and $n_{\text{t}}/n_{\text{r}}$ to refer to a specific transmit or receive antenna. We use the superscript $(\cdot)^{\text{RF}}$ to refer to the analog component of hybrid variables within the MIMO system.

1.6 Organization

The remainder of the dissertation is organized into three technical contributions described in Chapters 2-4 and concludes with directions for future work in Chapter 5. Chapter 2 introduces the beam management framework and beamspace learning for SSB codebook design. Chapter 3 extends the formulation to a hybrid, X-MIMO model and joint codebook learning for SSB and CSI-RS codebooks. Throughout these contributions, the learning was focused on beam training alignment. However, Chapter 4 reframes the learning process to maximize the achievable spectral efficiency. The strategy is presented in a multi-cell system to also show the network is capable of learning interference-minimizing beams. Conclusions and directions of future work are finalized in Chapter 5.

Chapter 2

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

2.1 Overview

In this chapter[1], we introduce the mathematical beam management model for a sub-6GHz, fully digital array. This mathematical model is an important contribution for building upon even as it is extended for hybrid and multi-cell scenarios in the upcoming chapters. In addition, we introduce beamspace transformations as a preprocessing step that enables effective codebook learning throughout this dissertation. The learning is setup as a one-step supervised learning paradigm. Finally, we use a novel feedback quantization method to efficiently evaluate the performance of AI/ML codebooks with

[1]This chapter is based upon my published work in (Dreifuerst and Heath 2023a): Dreifuerst, R. M. and Heath, R. W. (2023a). Machine learning codebook design for initial access and CSI type-II feedback in sub-6GHz 5G NR. *IEEE Trans. Wireless Commun.*.

various CSI configurations.

2.2 Introduction

Fifth-generation (5G) cellular systems have adopted a beam-based approach for MIMO communications. In such systems, broadcast control signals are beamformed to provide array gain (Dreifuerst and Heath 2023b). 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 (Lee et al. 2022).

Codebook-based beam management is a framework for beamforming and feedback that enables transmitting beamformed reference signals and multiple forms of CSI acquisition (Dreifuerst and Heath 2023b). 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 (Morozov et al. 2016). 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 (Raghavan et al. 2016). 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 (Heng et al. 2021; Giordani et al.

2019). To begin the process, the 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 (Alkhateeb et al. 2014). 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 narrow discrete Fourier transform (DFT) beamformers (Raghavan et al. 2016), 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 (Becirovic et al. 2022) or when the number of transmitting and receiving antennas on a UE is not equal (Samsung 2020).

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 proposal 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 evaluation 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 (Li et al. 2013), 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 (V. Va et al. 2018) 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 (Wang et al. 2021; Xue et al. 2022; Yang et al. 2022). These papers (Li et al. 2013; V. Va et al. 2018; Wang et al. 2021; Xue et al. 2022; Yang et al. 2022) 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 chapter, 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 (Wang et al. 2009; Alkhateeb et al. 2014; Xiao et al. 2016; Shafin et al. 2020; Xia et al. 2020; Xue et al. 2022; Heng et al. 2022; Dreifuerst et al. 2022). Hierarchical codebooks and the associated design process were proposed in (Wang et al. 2009; Alkhateeb et al. 2014; Xiao et al. 2016) 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 (Shafin et al. 2020), 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 (Xia et al. 2020), 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 (Xue et al. 2022) 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 (Heng et al. 2022), 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 (Heng et al. 2022) 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 (Dreifuerst et al. 2022). 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 (Morozov et al. 2016) 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 (Li et al. 2023), and as an indicator for beamforming vector selection (Li et al. 2021). 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 (Wen et al. 2018; Kim et al. 2021; Chen et al. 2022). All those works (Wen et al. 2018; Kim et al. 2021; Chen et al. 2022) 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 chapter, 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.

The remainder of this chapter 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

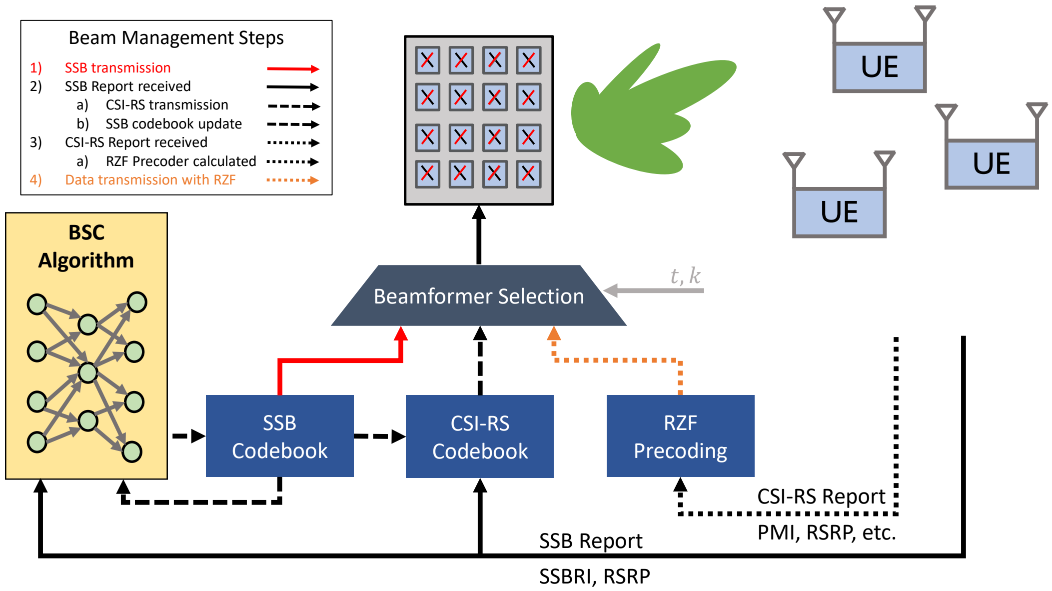

Figure 2.1: Beamforming codewords are used to transmit reference signals for synchronization from the SSB codebook. Based on the SSB report, additional reference signals are beamformed using the CSI-RS codebook for channel estimation by the UEs. Using the CSI-RS report, the BS can design a precoder that attempts to maximize the spectral efficiency. The proposed BSC algorithm updates the SSB codebook, which in turn affects the CSI-RS codebook and feedback.

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 this chapter. 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.

2.3 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 2.1. Throughout this chapter, 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.

2.3.1 MU-MIMO channel model

We model the system so that it is representative of real-world conditions to evaluate performance in a realistic beam training scenario for sub-6GHz massive MIMO. We do not explicitly specify TDD or FDD because this work does not rely on channel reciprocity. In addition, the channel will always be handled in the frequency domain as the multi-user downlink OFDM channel, however, we do not assume any specific channel model at this point. The OFDM MU-MIMO channel is defined for U users, over T time slots, and K subcarriers between each of the N_R receive antennas and N_T transmit antennas as $\mathbf{H} \in \mathbb{C}^{U \times T \times K \times N_\mathrm{R} \times N_\mathrm{T}}$. Note that we can generate this channel model using a clustered channel model as typically done by QuaDRiGa (Jaeckel et al. 2014), 3GPP, and other raytracing or statistical models with the array response vectors $\mathbf{a}_N(\theta) = \frac{1}{\sqrt{N}}[1, e^{j\pi\cos\theta}, e^{j2\pi\cos\theta}, ..., \ e^{j(N-1)\pi\cos\theta}]^T$. Defining the time and frequency responses for every antenna pair is critical to accurately model the beam training process with specific resource elements allocated for pilots (reference signals) and data. With this setup in mind, a generic received signal model for signal $\mathbf{s}$ transmitted with precoder/beamformer $\mathbf{F}$ and received with a combiner $\mathbf{W}$ over a band of subcarriers K is

$$\mathbf{D}_{t,k} = \sum_{p=0}^{U-1} \mathbf{F}_{p,t,k}\mathbf{s}_{p,t,k} \tag{2.1}$$

$$\mathbf{y}_{u,t,k} = \frac{1}{\sqrt{KN_{\mathrm{T}}}}\mathbf{W}_{u,t,k}\left(\mathbf{H}_{u,t,k}\mathbf{D}_{t,k} + \mathbf{N}_{u,t,k}\right). \tag{2.2}$$

The noise $\mathbf{N}$ has entries with zero mean and variance according to the thermal noise and the noise figure of the receiver. During the next subsections, we will specify how $\mathbf{W}_{(u)}, \mathbf{F}_u$, and $\mathbf{s}_u$ are used during the SSB, CSI-RS, and data transmission stages. We will assume throughout the chapter that all beamformers are normalized according to a per-symbol power constraint, i.e. $\|\mathbf{F}\|^2 = \ N_{\mathrm{T}}$. In the SSB and CSI-RS stages, $\mathbf{F}$ is selected from one or more codebooks that play critical roles in beam management.

2.3.2 Codebooks

Communication protocols employ codebooks in many different ways for beam training, feedback representations, information quantization, and more. Generally, codebooks are understood as an ordered set of elements called codewords. Throughout this work on beam management, codewords will correspond to a vector of complex values and we will treat codebooks as matrices to simplify notation in later steps. There are three codebooks that we will reference and directly impact beam management: a codebook of beamforming vectors for SSB transmission (SSB codebook), a codebook for beamforming CSI-RS transmissions (CSI-RS codebook), and a codebook for feedback quantization (FB codebook). The sizes of the codebooks are based on the supported usage for each one. The SSB codebook $\mathbf{B}^{\text{SSB}} \in \mathbb{C}^{N_\text{P} \times L_{\max}}$ has $L_{\max}$ codewords which is between 1 and 64 depending on the carrier frequency (Dreifuerst and Heath 2023b). The number of logical ports, N_P, determines the digital precoding dimension and is equal to $N_\text{X}N_\text{Y}$ in a fully digital array. The CSI-RS codebook $\mathbf{B}^{\text{CSI-RS}} \in \mathbb{C}^{N_\text{P} \times N}$ is at 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 $\mathbf{B}^{\text{DFT}} \in \mathbb{C}^{N_\text{P} \times N_\text{X} O_\text{H} N_\text{Y} O_\text{V}}$ 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 interrelationships 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 2.5b, 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.

2.3.3 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 (Giordani et al. 2019), 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) (Giordani et al. 2019). These SSBs are beamformed using a specific beamforming vector selected from $\mathbf{B}^{\mathrm{SSB}}$ 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 $L_{\max} = \{1, 4, 8, 64\}$ 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 K^{SSB} frequency resources at times $T_i^{\mathrm{SSB}} i$ for the i^{th} SSB using beamformer $\mathbf{f}_i \in \mathbf{B}^{\mathrm{SSB}}$. 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 $\mathbf{B}^{\mathrm{SSB}}$. During SSB reception, the generic received signal model from (2.2) is modified to account for the single data stream

$$y_{u,t,k}^{\mathrm{SSB}_i} = \frac{1}{\sqrt{KN_{\mathrm{T}}}}\mathbf{w}_{u,t,k}^{*}\left(\mathbf{H}_{u,t,k}\mathbf{f}_i s_{t,k} + \mathbf{n}_{u,t,k}\right). \tag{2.3}$$

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 (3GPP 2023a). With the received signal model (2.3) in mind, we now define an important quantity provided as feedback during initial access, the RSRP.

2.3.4 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 $n_{\mathrm{r}} \in \{0, ..N_{\mathrm{R}}-1\}$ 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

$$\mathbf{y}_{u,t,k}^{\mathrm{SSB}_i} = \frac{1}{\sqrt{KN_{\mathrm{T}}}}\mathbf{H}_{u,t,k}\mathbf{f}_i s_{t,k} + \mathbf{n}_{u,t,k} \tag{2.4}$$

$$\mathrm{RSRP}_{i,u} = \max_{n_{\mathrm{r}}} \sum_{k\in\mathbf{K}_{\mathrm{DMRS}}} \sum_{t\in\mathbf{T}_{i,\mathrm{DMRS}}} \left\|\mathbf{y}_{u,t,k}^{\mathrm{SSB}_i}\right\|^2. \tag{2.5}$$

Note that (2.5) 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

$$\mathbf{p} = \left\{ \max_i \mathrm{RSRP}_{i,u} \right\}_{u=0}^{U-1} \tag{2.6}$$

$$\mathbf{m} = \left\{ \arg\max_i \mathrm{RSRP}_{i,u} \right\}_{u=0}^{U-1}. \tag{2.7}$$

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.

2.3.5 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 (Giordani et al. 2019; Heng et al. 2021; Samsung 2020) 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 (Morozov et al. 2016; Xue et al. 2022; Singh and Joshi 2021) (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 $N_{\text{X}}N_{\text{Y}}O_{\text{H}}O_{\text{V}}$ oversampled DFT vectors throughout this paper. The selection of the CSI-RS beams from SSB decomposition is further refined in Section 2.4.

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 $(\mathbf{T}^{\mathrm{CSI}_i}, \mathbf{K}^{\mathrm{CSI}_i})$, 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 (Morozov et al. 2016) and references therein. At this stage, the entire UE antenna array is active and the UE can perform maximum ratio combining $\mathbf{w}^i_{u,t,k} = \left(\widehat{\mathbf{H}_{u,t,k}\mathbf{f}_i}\right)^*$ based on the estimated beamformed channel and appropriately normalized. The estimated SNR at the receiver is

$$\mathrm{SNR}_u = \max_{c \in P_{\mathrm{CSI}}} \frac{1}{K} \sum_{k \in \mathbf{K}_{\mathrm{CSI}_c}} \sum_{t \in \mathbf{T}_{\mathrm{CSI}_c}} \frac{1}{N_{\mathrm{T}}} \frac{\left\| \widehat{\mathbf{H}_{u,t,k}\mathbf{f}_c} \right\|^2}{\mathbb{E}[\|\mathbf{N}_{u,t,k}\|^2]}. \tag{2.8}$$

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 $(\mathbf{H}_{u,t,k}\mathbf{F}^{\text{CSI-RS}})$.

2.3.6 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 (Heath Jr. and Lozano 2018, section 3.7)

$$\text{MSE} = \frac{1}{\frac{N_{\text{pilots}}}{N_T}\text{SNR}}. \tag{2.9}$$

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 (2.9). 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, $\widehat{\mathbf{HF}}$, 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 $\widehat{\mathbf{HF}}$). The BS could then use this information to beamform to a given user and null steer (Godara 1997) the beams for other users. This puts the majority of the computational burden, however, on the UE. Instead, we follow models similar to (Castellanos et al. 2018) 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.

2.3.7 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 (3GPP 2023a), 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 2.5. 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 (MathWorks 2023). 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. (Heath et al. 2016). Then we can define a diagonal oversampling matrix $\mathbf{D}_{O,N}$ for an oversampling O and N elements as

$$\mathbf{D}_{O,N} = \mathrm{diag}\left(1,\ e^{\left(-\frac{\mathrm{j}2\pi}{ON}\right)}, ..., e^{\left(-\frac{\mathrm{j}2\pi(N-1)}{ON}\right)}\right). \tag{2.10}$$

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 $\mathbf{D}_{O,N}$

$$\mathbf{B}_{O,N} = [\mathbf{D}^{0}_{O,N}\mathbf{U}_N, \mathbf{D}_{O,N}\mathbf{U}_N, ..., \mathbf{D}^{O-1}_{O,N}\mathbf{U}_N]. \tag{2.11}$$

Finally, the UPA oversampling codebook can be defined as the kronecker product of the N_{X} and N_{Y} dimensions

$$\mathbf{B}^{\mathrm{DFT}} = \mathbf{B}_{O_{\mathrm{H}},N_{\mathrm{X}}} \otimes \mathbf{B}_{O_{\mathrm{V}},N_{\mathrm{Y}}} \in \mathbb{C}^{N_{\mathrm{X}}N_{\mathrm{Y}}\times O_{\mathrm{H}}N_{\mathrm{X}}O_{\mathrm{V}}N_{\mathrm{Y}}}. \tag{2.12}$$

With the FB codebook defined in (2.12), the PMI selection procedure can begin. The UE first calculates the inner product for all beams from the oversampled codebook $\mathbf{B}^{\mathrm{DFT}}$ and selects the first beam (and by association the oversampling basis) as

$$\mathbf{C} = \widehat{\mathbf{HF}}\mathbf{B}^{\mathrm{DFT}} \ \in \mathbb{C}^{N_{\mathrm{R}}\times N_{\mathrm{X}}O_{\mathrm{H}}N_{\mathrm{Y}}O_{\mathrm{V}}} \tag{2.13}$$

$$Q_0 = \arg\max|\mathbf{C}| \tag{2.14}$$

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 $\mathbf{B}^{\mathrm{DFT}}$, 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

$$\mathbf{B}_{\text{ortho},Q_0} = \mathbf{D}_{O_{\text{H}},N_{\text{X}}}^{Q_{0,X}} \mathbf{U}_{N_{\text{X}}} \otimes \mathbf{D}_{O_{\text{V}},N_{\text{Y}}}^{Q_{0,Y}} \mathbf{U}_{N_{\text{Y}}} \tag{2.15}$$

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

$$\mathbf{c}_{\text{ortho}} = \widehat{\mathbf{HF}}_{n_{\text{r}}} \mathbf{B}_{\text{ortho},Q_0} \ \in \mathbb{C}^{N_{\text{X}} N_{\text{Y}}} \tag{2.16}$$

$$\mathbf{Q}_{n_{\text{r}}} = (L_{\text{CSI}})\text{- arg max} \, |\mathbf{c}_{\text{ortho}}| \quad \forall \ n_{\text{r}} \tag{2.17}$$

$$\mathbf{A}_{n_{\text{r}}} = \mathbf{c}_{\text{ortho}} \left[\mathbf{Q}_{n_{\text{r}}} \right] \quad \forall \ n_{\text{r}} \tag{2.18}$$

$$\mathbf{A}_{n_{\text{r}}} = \frac{\mathbf{A}_{n_{\text{r}}}}{\mathbf{A}[n_{\text{r}}, \arg\max |\mathbf{A}_{n_{\text{r}}}|]} \quad \forall \ n_{\text{r}}. \tag{2.19}$$

First, the orthogonal beams are "evaluated" in (2.16), then we make use of the k- arg max operation defined as the k indices for the k largest values in (2.17) to select the L_{CSI} beam indices. Finally, the combining coefficients are calculated as the complex values of the beam products (2.18) and normalized according to the largest magnitude coefficient (2.19). 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 $\binom{N_{\mathrm{X}}N_{\mathrm{Y}}}{L_{\mathrm{CSI}}}$ combinations and one index corresponding to the oversampling factors of the strongest index $(Q_{0,X}, Q_{0,Y})$ is used. In this chapter, 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 $\mathrm{S_B} = 8$ in FR1, to mitigate the growing feedback and complexity. All of the previous and subsequent steps are defined for $\mathrm{S_B} = 1$, 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 $(L_{\mathrm{CSI}}, \mathrm{S_B})$ impacts performance in Section 2.5.

2.3.8 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

$$\widehat{\widehat{\mathbf{HF}}}[n_\mathrm{r}] = \sum_{\ell=0}^{L_\mathrm{CSI}-1} \mathbf{A}[n_\mathrm{r}, \ell]\mathbf{B}^\mathrm{DFT}[:, \mathbf{Q}_{n_\mathrm{r},\ell}] \quad \forall\ n_\mathrm{r}. \tag{2.20}$$

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

$$\widehat{\widehat{\mathbf{H}}}_u = \widehat{\widehat{\mathbf{H}_u\mathbf{F}}}\mathbf{F}^\dagger_\text{CSI-RS}. \tag{2.21}$$

In the case of DFT beams, this process simply becomes an IDFT operation (Morozov et al. 2016), which is significantly more efficient than the matrix inverse of an $N_\mathrm{P} \times N_\mathrm{P}$ 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 $\widehat{\widehat{\mathbf{H}}} \in \mathbb{C}^{U\times N_\mathrm{R}\times N_\mathrm{P}}$

$$\widehat{\widehat{\mathbf{H}}} = \left[\widehat{\widehat{\mathbf{H}}}_0, \widehat{\widehat{\mathbf{H}}}_1, ..., \widehat{\widehat{\mathbf{H}}}_{U-1}\right]. \tag{2.22}$$

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 (Heath Jr. and Lozano 2018, section 9.9). The RZF precoder for each user is

$$\mathbf{F}_u = \frac{(\sum_{i=0}^{U-1} \widehat{\mathbf{H}}_i^* \widehat{\mathbf{H}}_i + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^2])^{-1}\widehat{\mathbf{H}}_u^*}{\left\|(\sum_{i=0}^{U-1} \widehat{\mathbf{H}}_i^* \widehat{\mathbf{H}}_i + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^2])^{-1}\widehat{\mathbf{H}}_u^*\right\|}. \tag{2.23}$$

It is assumed all precoders are appropriately normalized to $\sqrt{\frac{N_{\mathrm{T}}}{U}}$ 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

$$\mathbf{I}_{u,t,k} = \left(\sum_{i=0}^{U-1} \mathbf{H}_{u,t,k}\mathbf{F}_i\mathbf{F}_i^*\mathbf{H}_{u,t,k}^* + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^2]\right)^{-1} \tag{2.24}$$

$$\mathrm{SINR}_{u,t,k,r} = \frac{1}{K}\frac{\mathbf{F}_u^*[:,r]\mathbf{H}_{u,t,k}^*\mathbf{I}_{u,t,k}\mathbf{H}_{u,t,k}\mathbf{F}_u[:,r]}{1-\mathbf{F}_u^*[:,r]\mathbf{H}_{u,t,k}^*\mathbf{I}_{u,t,k}\mathbf{H}_{u,t,k}\mathbf{F}_u[:,r]}. \tag{2.25}$$

While the SINR expression (2.25) 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, $\mathrm{SE}_{u,t}$, is

$$\mathrm{SE}_{u,t} = \sum_{k=0}^{K-1}\sum_{r=0}^{R-1} \log_2(1 + \mathrm{SINR}_{u,t,k,r}). \tag{2.26}$$

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 $(T_{\mathrm{BM}}, K_{\mathrm{BM}})$ from the spectral efficiency calculation

$$\mathrm{ESSE} = \sum_{u=0}^{U-1} \sum_{\substack{t \notin T_{\mathrm{BM}} \\ k \notin K_{\mathrm{BM}}}} \sum_{r=0}^{R-1} \log_2(1 + \mathrm{SINR}_{u,t,k,r}). \tag{2.27}$$

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

2.4 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 SSB codebook. The full methodology–processing, architecture, training, and evaluation–is critical to integrating wireless domain knowledge and physical structure

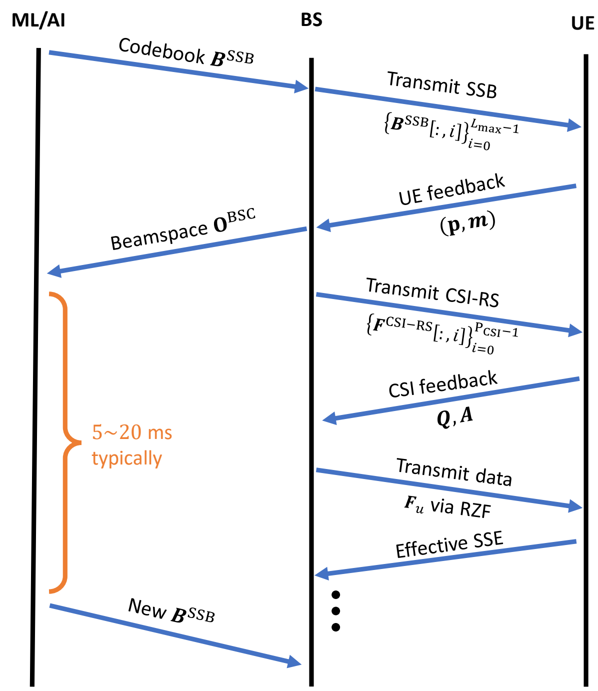

Figure 2.2: A timing diagram of the process for codebook learning and evaluation. The AI/ML engine provides new codebooks at each SSB interval for the BS. At each CSI-RS period, the BS transmits pilots using the CSI-RS codebook, which is determined based on the SSB codebook and feedback. The CSI-RS feedback is used to determine the precoder used for data transmission.

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

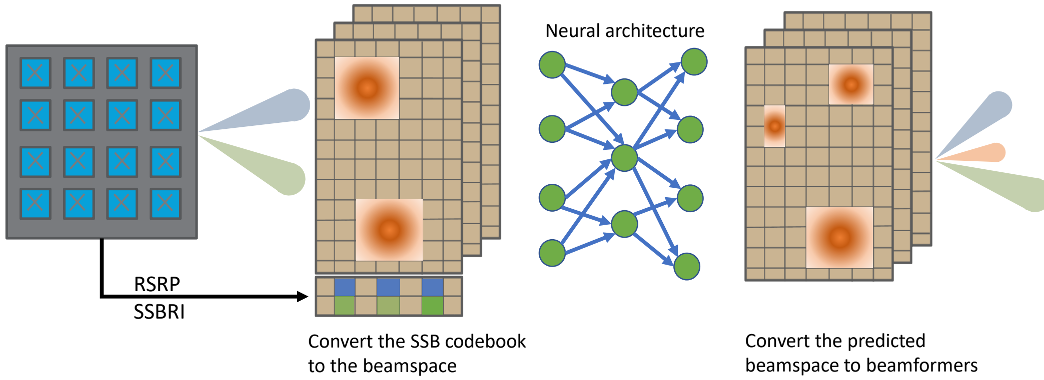

Figure 2.3: A depiction of the BSC pre-processing, neural architecture, and post-processing. The feedback is converted to an angular representation (beamspace) and the strongest directions are identified based on the feedback reported. The neural network jointly processes the SSB feedback and beamspace 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.

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 (Sayeed 2002) that has been shown to be beneficial for learning in similar settings (Wen et al. 2018). This ensures that regardless of the specific beamforming 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

$$\boldsymbol{\theta}_{N_m} = \frac{1}{\pi}[0, 1, ..., N_m - 1]^T \tag{2.28}$$

$$\mathbf{U}_{N_s,N_m} \triangleq [\mathbf{a}_{N_s}(\boldsymbol{\theta}_0), ...\mathbf{a}_{N_s}(\boldsymbol{\theta}_{N_m-1})]. \tag{2.29}$$

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

$$\mathbf{F}_i^{\mathrm{SSB}} = [\mathbf{B}^{\mathrm{SSB}}_{0:N_{\mathrm{X}},i}, ..., \mathbf{B}^{\mathrm{SSB}}_{(N_{\mathrm{Y}}-1)N_{\mathrm{X}}:N_{\mathrm{X}}N_{\mathrm{Y}},i}] \tag{2.30}$$

$$\mathbf{O}_i^{\mathrm{BSC}} = \mathbf{U}^{*}_{N_{\mathrm{X}},N_{x0}} \mathbf{F}_i^{\mathrm{SSB}} \mathbf{U}_{N_{\mathrm{Y}},N_{y0}} \quad \forall i. \tag{2.31}$$

The codebook must be reshaped from a vector of size N_{T} to the planar dimensions $N_X \times N_Y$ in (2.30) before the beamspace conversion in (2.31) that produces the beamspace observation, $\mathbf{O}^{\mathrm{BSC}}$. 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

$$\mathbf{O}^{\mathrm{BSC}} = \{[\mathbf{O}^{\mathrm{BSC}}[i], \sum \mathbf{1}_{\mathbf{m}=i}, \mathbf{1}^T_{\mathbf{m}=i}\mathbf{p}]\}_{i=0}^{L_{\max}-1}. \tag{2.32}$$

With the pre/post-processing defined, we can now explain 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 2.5 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 (Raghavan et al. 2016). 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 (Ho et al. 2020) networks), and translation-based architectures (transformers (Liu et al. 2021) 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 2.1. A comparison of model architectures is shown in Figure 2.7 within Section 2.6. The model is trained using an Adam optimizer (Kingma and Ba 2014) with early stopping and cosine learning rate decay. Training is performed with a $550,000$ sample training dataset following the setup in Section 2.6 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 ($d_{\cos}(\mathbf{a},\mathbf{b}) = 1 - \mathbf{a}^*\mathbf{b}/(\|\mathbf{a}\|\,\|\mathbf{b}\|)$ (Chen et al. 2022))) 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 (> 5dB) performance due to focusing on the pixel-wise magnitude. In comparison, the cosine distance focuses on the

Table 2.1: Beamspace-Codex network parameters

Layer	**Primary Parameter**	**Act.**	**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			$N_{x0} \times N_{y0} \times 2L_{\max}$

relative magnitude or direction of the RSV beamspace.

2.5 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 (Jaeckel et al. 2014) 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 $N_\mathrm{X} = 4$, $N_\mathrm{Y} = 8$ planar array using a 3GPP 3D antenna model with half-wavelength spacing. Each user has $N_\mathrm{R} = 4$ antennas all tuned for a 3.5GHz carrier frequency as might be found in modern UEs which are typically capable of 4 layer MIMO (Yang et al. 2019). 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 $[0, 3]$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.

2.6 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 (Dreifuerst and Heath 2022), 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.

2.6.1 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 2.4a 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 $N_{\mathrm{T}} = 16$ 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 2.4b, 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.

2.6.2 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 (3GPP 2023a). User selection during the data phase is performed based only on the knowledge at the BS by first estimating the sum spectral efficiency achieved with the imperfect CSI for all users. The best estimated SE combination is selected and used until

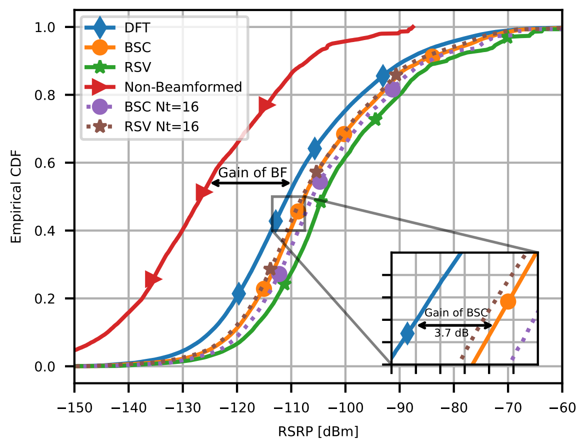

(a) SSB RSRP

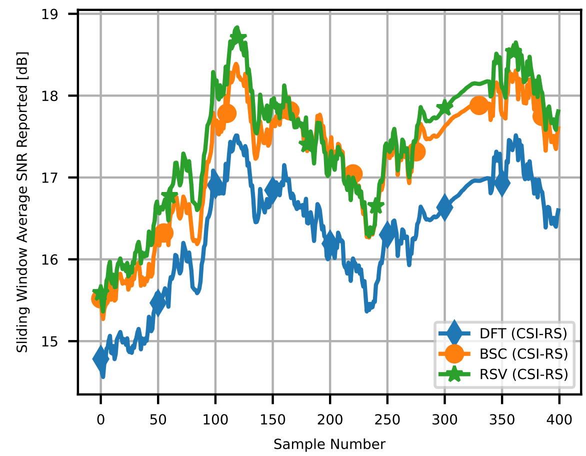

(b) CSI-RS SNR

Figure 2.4: 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.

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 2.5a. 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. 2.5b 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 2.5a-2.6a, a good balance of feedback and performance can be achieved with a setting of $L_{\text{CSI}} = 32$, $\text{S}_\text{B} = 1$, 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 $L_{\text{CSI}} <= 6$, 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 perfor-

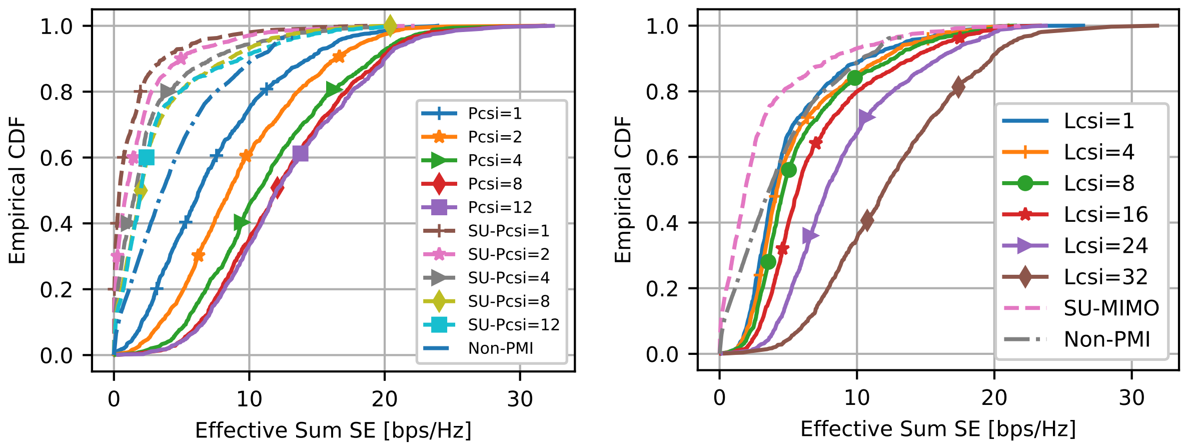

(a) Codebook size P_{CSI} (b) Quantization paths L_{CSI}

Figure 2.5: 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 $L_{\mathrm{CSI}} = 32$ for (a), $P_{\mathrm{CSI}} = 16$ for (b), $\mathrm{S_B} = 1$, $\mathrm{NRB} = 24$.

mance. 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.

2.6.3 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

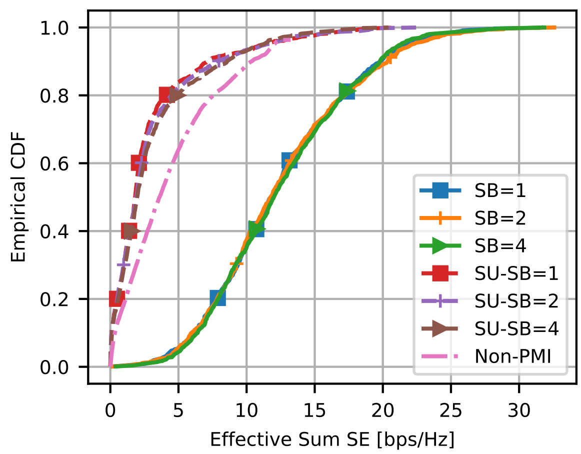

(a) Subbands S_B

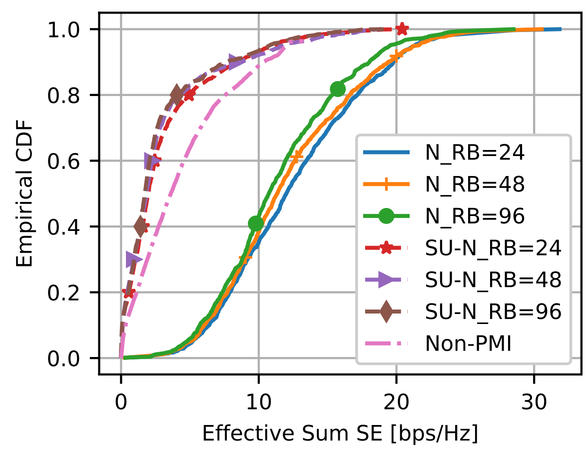

(b) Resource blocks per subband N_{RB}

Figure 2.6: 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 $L_{CSI} = 32$, $P_{CSI} = 16$, 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.

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 necessary. Figure 2.7 shows the proposed BSC network outperforming different architectures trained with the same data, optimization strategy, and comparable execution time by about 1dB on average.

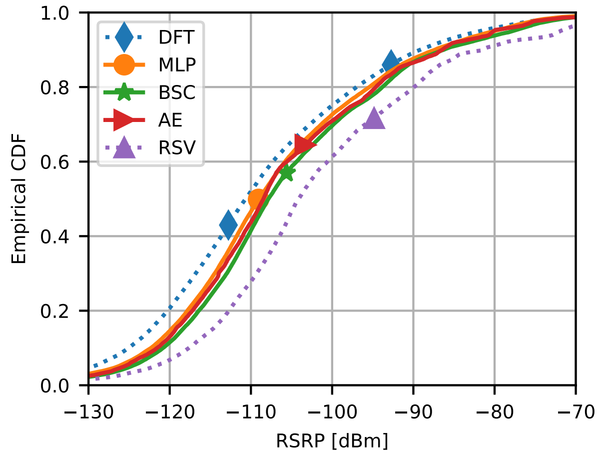

Figure 2.7: A comparison of the RSRP performance of classic neural architectures–multilayer perception (MLP), convolutional autoencoder (AE)–within the beamspace framework. It can be seen that although the architectures have comparable results, the BSC architecture based on CNNs is the best choice for learning beamspace features.

2.6.4 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 2.8 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 2.8 shows the performance in the new environment (agnostic) is noticeably worse relative to

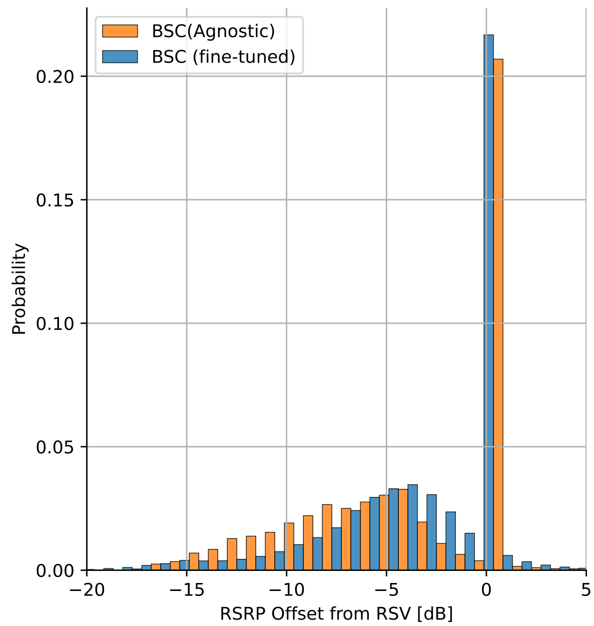

Figure 2.8: Histogram representing the performance delta between BSC and RSV codebooks in the agnostic setting and after fine-tuning. It can be seen that the worst-case differences are resolved and in some cases (when more users are present than SSB beams) the RSV codebook can be outperformed based on the higher proportion of the blue columns to the right of the orange columns.

the RSV. The performance loss is improved during fine-tuning, especially for the lowest 10% of users.

2.7 Conclusion

In this chapter, we presented a framework for learning dynamic SSB codebooks. We set up a realistic beam management scenario in sub-6GHz 5G NR with multiple UEs and multi-stream communication and fully digital arrays at both sides of the wireless link. We used this framework to determine that machine learning can design SSB codebooks more effectively than traditional codebooks. The proposed Beamspace-Codex algorithm more effectively learns codebooks than direct implementations by first transforming codebooks into the beamspace. This representation is shown to generalize effectively and achieve strong initial access reception than traditional codebooks. We also investigated how feedback configuration impacts performance. This showed that CSI quantization resolution was a limiting factor in MU-MIMO communications due to the multipath channel introducing interference. Using our channel decomposition strategy, higher resolution feedback is not computationally restrictive. These results highlighted the need for beamformed channel measurement as an initial step for reducing the amount of CSI needed to accurately represent CSI.

Chapter 3

Hierarchical ML Codebook Design for Extreme MIMO Beam Management

3.1 Overview

In this chapter[1], we build upon the previous chapter to learn joint SSB and CSI-RS beamforming codebooks. We employ the same beamspace preprocessing in a parallel architecture to generate the codebooks for extreme MIMO systems. Then, we propose a simple correlation-based beam selection procedure to calculate the active CSI-RS beam subset that supports beam refinement. We also show that beamspace learning is capable of generalizing across a range of scenarios that validate the learning algorithm.

[1] This chapter is based upon my work in (Dreifuerst and Heath 2024a): Dreifuerst, R. M. and Heath, R. W. (2024a). Hierarchical ML codebook design for extreme MIMO beam management. *Submitted to IEEE Transactions on Machine Learning in Communications and Networking.*

3.2 Introduction

Large-scale or Massive MIMO (M-MIMO) has been an important technology in mobile broadband since the first investigations on 3D beamforming in 3GPP release 13 (Dreifuerst and Heath 2023b). Since then, larger arrays have been envisioned with extremely large arrays ($\gg 64$ antennas) for 6G (Holma et al. 2021). Effectively employing M-MIMO or X-MIMO, however, requires obtaining accurate CSI. With the inception of 5G, M-MIMO was supported through a process called beam management which integrated beamforming for control signaling, hierarchical codebooks for beam training, and enabled multiple configurations of feedback (Dreifuerst and Heath 2023b). Notably, beam management supports hybrid arrays where the number of antenna ports visible to digital communication systems and the number of antennas in the array might be vastly different, as found in the so-called hybrid MIMO architecture (Heath et al. 2016). While the beam management framework is designed to be flexible, there has been limited research on optimizing beam management codebooks for this flexible framework.

Beam management is a multi-step process that includes transmitting beamformed reference signals and obtaining multiple forms of CSI to balance the overhead, latency, and accuracy (Heng et al. 2021; Giordani et al. 2019; Dreifuerst and Heath 2023b). In 5G, the first step is synchronization which requires transmitting one or more synchronization signal blocks with a beamformer from a codebook (denoted as the *SSB codebook*). These are typically understood as wide-area coverage beams because synchronization does not require a high signal-to-noise ratio, but wide-area beamforming improves cell coverage and aids the subsequent beam search. From the set of SSBs, user equipment provides a small feedback message identifying the strongest SSB and the power received from the transmission (Giordani et al. 2019). Based on the SSB feedback, the base station selects a refined beamforming codebook (identified as the *CSI-RS codebook*) for transmitting CSI reference signals, which contain pilot signals for channel estimation. These two beamforming codebooks support different functions (initial access and synchronization compared to beam refinement and channel estimation) and operate over different time-frequency resource allocations to support these roles. Following CSI-RS reception, UEs estimate the effective channel and quantize it according to another codebook (the *feedback codebook*) based on the BS array geometry. Finally, the CSI feedback is provided to the BS and data transmission can begin with an appropriately designed precoder to serve one or more users.

Optimizing the various codebooks is challenging due to the difficulty in mathematically capturing the relationships between the channel, mobile users, and the codebook interrelationships. One possible solution is applying AI/ML to learn these underlying characteristics from data. While AI/ML has become an integral part of domains like computer vision and language modeling, mobile broadband has not seen as many benefits due to the well-known structure and strict timing requirements of wireless communications (Wang et al. 2017; Zhang et al. 2019). Recently, 3GPP has announced investigations for the integration of AI/ML targeting 3GPP Release 18 and beyond (Lin 2023). In particular, beam management and CSI feedback are identified as key areas where learning strategies are envisioned (Lin 2023).

With the initial proposals for 6G underway, new research on AI/ML for extremely large arrays, equipped with hundreds of antennas, is a critical topic, especially operating in new communication bands (Holma et al. 2021). 5G initially began supporting massive MIMO in a hybrid format, where analog beam training was performed using the SSB and CSI-RS beamformers, and digital precoding was based on the CSI feedback. Given the support and success of these formats, it is valuable to design and optimize strategies for X-MIMO within this formulation. Successfully accomplishing this task paves the way for beam management with AI/ML in 6G.

There are three directions of work related to AI/ML beam management: 1) beam training, 2) codebook design, and 3) CSI feedback. First, we review a set of relevant prior work incorporating machine learning and beam training. Beam training has received the most attention of the three directions when it comes to machine learning applications. Classical data-driven techniques were explored in (Li et al. 2013; V. Va et al. 2018) using historical beam training data. While the codebooks and algorithms rely on more traditional techniques, the concept of "learning" a reduced codebook from site-specific data is a foundational concept for more recent deep learning techniques (Wang et al. 2021; Xue et al. 2022; Yang et al. 2022, 2023). These papers ((Wang et al. 2021; Xue et al. 2022; Yang et al. 2022, 2023)) focus on millimeter-wave (mmWave) channels and typically exploit the sparsity to aid the beam training process. Furthermore, the codebooks within these papers are standard codebooks like DFT which are unlikely to perform as well in rich scattering environments as seen in upper-mid bands (7-20GHz) envisioned for 6G (Holma et al. 2021). In addition, that work (Wang et al. 2021; Xue et al. 2022; Yang et al. 2022, 2023) does not consider a MIMO format with multiple UE antennas and streams of data, instead limiting the setting to analog beamforming. In this chapter, we consider a scenario with multiple users, multiple data streams, and frequency selective, OFDM channels in rich multipath environments.

Codebook and beamformer design using machine learning has received some focus in recent years (Shafin et al. 2020; Xia et al. 2020; Xue et al. 2022; Heng et al. 2022; Dreifuerst et al. 2022). Hierarchical codebooks, which are a core concept in 5G, were proposed in (Wang et al. 2009; Alkhateeb et al. 2014; Xiao et al. 2016) to separate the beam search process into multiple resolutions to reduce the overhead of beam search. Deep reinforcement learning for broadcast beam pattern transmission was shown to be effective in small-scale MIMO (Shafin et al. 2020). Supervised learning for SSB codebook design in narrowband MIMO was also shown to improve over standard codebooks (Dreifuerst et al. 2022). Unsupervised learning was applied to beamformer design from narrowband channel measurements (Xia et al. 2020) but required perfect channel state information. Broadcast probing beam patterns were learned using channel state information in (Heng et al. 2022) which is one of the only papers to consider a hierarchical strategy similar to 5G based on SSB and CSI-RS codebooks. This work (Heng et al. 2022) provides motivation for researching realistic deep learning codebook methods, but it remains unclear how the integration of ML impacts the overall system performance as opposed to simply improving the beam alignment. Our work addresses this question while further increasing the role that data-driven algorithms play in codebook operations.

While codebook design is typically focused on improving the received signal power, codebooks are also employed for quantizing channel state information into discrete feedback. ML has also been applied to improve feedback quantization in recent work (Wen et al. 2018; Kim et al. 2021; Chen et al. 2022; Xiao et al. 2023). In general, ML methods propose encoder-decoder structures (Wen et al. 2018; Kim et al. 2021; Chen et al. 2022; Xiao et al. 2023), that employ jointly-trained models at the UE and BS. This type of application involves significant overhead, however, to ensure both terminals have synchronized neural network weights or parameters. Furthermore, the system effects in MU-MIMO settings are not considered, even though high resolution CSI feedback is primarily designed for improving MU-MIMO performance (Morozov et al. 2016). Throughout these investigations, CSI optimization is treated as a separate task beam training and channel estimation, thereby neglecting the relationships throughout the system. In this chapter, we incorporate beamformed channel feedback as a result of beam management to study the effects of ML codebooks on the overall system performance.

In our prior work (Dreifuerst and Heath 2023a), we focused on a 5G, fully digital sub-6GHz system. From that investigation, we found that significant gains can be achieved with improvements to the beam management framework using deep learning. While the strategy was effective for SSB transmission, we now extend the consideration to a new system configuration that supports hybrid, upper-mid band systems with jointly-learned SSB and CSI-RS codebooks. Throughout the process, the overall goal is to balance the computational complexity, cost, and performance of CSI acquisition to maximize system-level performance. Our contributions in this manuscript are as follows:

- First, we precisely define the beam management process for an envisioned hybrid, X-MIMO system. The process includes all relevant steps from initial access up through achievable network spectral efficiency calculation. This system enables a holistic evaluation of codebooks that captures the interdependencies of a realistic beam management strategy for 6G and is a natural evolution of that supported in 5G.

- Second, we design a multi-task neural network architecture and processing pipeline to design SSB and CSI-RS codebooks. The pipeline starts by converting the current SSB codebook and SSB feedback into the beamspace. Then the input is fed into two neural networks that are jointly trained in an end-to-end fashion to output an SSB and a set of candidate CSI-RS codebooks. This architecture is formulated to fit the strict timing requirements of beam training. The joint neural network formulation is able to learn codebooks that achieve beam training results similar to if perfect CSI were known a priori.

- Third, we rigorously evaluate the capabilities of the beam management framework in X-MIMO using various codebooks, physically array geometries, and feedback strategies. The most important result we find is that the proposed system, with limited feedback, can achieve comparable performance to DFT-codebook systems with full feedback in MU-MIMO settings. Furthermore, we find that the X-BM design transfers efficiently between new environments, antenna arrays, and frequency ranges.

The remainder of this chapter is organized as follows. First, we introduce the system model and briefly highlight the SSB and CSI-RS beam management processes. We conclude this section with an overview of PMI and define the data transmission stage. Next, we introduce the proposed algorithm, X-BM, that includes beamspace preprocessing, a multi-stage neural network, and end-to-end learning. In the final sections, we present the simulation setup and a rigorous evaluation of the various codebooks, feedback formats, and overall performance achieved by our proposed method.

3.3 System model

We consider one region served by a single X-MIMO base station. The BS is equipped with an $N_{\mathrm{T}}^{\mathrm{RF}} = N_{\mathrm{X}}^{\mathrm{RF}} \times N_{\mathrm{Y}}^{\mathrm{RF}}$ planar array in a hybrid architecture such that antenna elements are connected with phase shifters to radio frequency chains to form a digital dimension of size $N_{\mathrm{T}} = N_{\mathrm{X}} \times N_{\mathrm{Y}}$. The phase shifters are limited to b_{phase} bits resolution but we assume a fully-connected structure such that every digital port is connected to each physical antenna with a separate phase shifter (Alouzi et al. 2023). On the other side of the radio link, there are U UEs, where U is a random variable, each equipped with an N_{R} uniform linear array which we assume to be fully digital for this investigation. This reduces the notational complexity and allows us to focus on the BS operations, whereas UE-side optimization is left to future work. All antennas are half-wavelength spaced in the YZ planes. It remains to be seen what bands and configurations will be standardized in 6G, so we will assume channels are generated according to a model but rely on machine learning to capture channel structures rather than incorporate a specific channel model into the system. The OFDM MU-MIMO channel is defined for U users, over T time slots, and K subcarriers between each of the terminals as

$$\mathbf{H} \in \mathbb{C}^{U \times T \times K \times N_{\mathrm{R}} \times N_{\mathrm{T}}^{\mathrm{RF}}}. \tag{3.1}$$

Note that we can generate this channel model using a clustered channel model as typically done by Sionna (Hoydis et al. 2022), 3GPP, and other raytracing or statistical models and array response vectors $\mathbf{a}_N(\theta, \phi) = \frac{1}{\sqrt{N}}[1, e^{\mathrm{j}\pi \cos\theta \sin\phi}, e^{\mathrm{j}2\pi \cos\theta \sin\phi}, ..., \ e^{\mathrm{j}(N-1)\pi \cos\theta \sin\phi}]^T$ for the path direction (θ, ϕ).

With this setup in mind, we work with the following generic received signal model for signals $\mathbf{s}$ transmitted with digital and analog precoders $\mathbf{F}_{u,t,k}$, $\mathbf{F}_{c,t}^{\mathrm{RF}}$ and received with combiners $\mathbf{W}_{u,t,k}$

$$\mathbf{y}_{u,t,k} = \frac{1}{\sqrt{N_{\mathrm{T}}}} \mathbf{W}_{u,t,k}^* \mathbf{H}_{u,t,k} \mathbf{F}_{c,t}^{\mathrm{RF}} \sum_{\mathrm{u}=0}^{U-1} \mathbf{F}_{\mathrm{u},t,k} \mathbf{s}_{\mathrm{u},t,k} + \mathbf{W}_{u,t,k}^* \mathbf{N}_{t,k}. \tag{3.2}$$

The noise $\mathbf{N}$ is modeled as independent complex Gaussian random values to account for thermal noise and the noise figure of the receiver. Our model assumes that timing and frequency synchronization have been performed, which is supported through the embedded primary synchronization and secondary synchronization sequences contained within the SSB signals. During the next subsections, we will specify how $\mathbf{W}, \mathbf{F}$, and $\mathbf{s}$ are determined during the SSB, CSI-RS, and data transmission stages according to an SSB codebook

$\mathbf{B}^{\text{SSB}} \in \mathbb{C}$, a CSI-RS codebook $\mathbf{B}^{\text{CSI-RS}} \in \mathbb{C}$, and a predefined feedback codebook. We will assume throughout the chapter that all digital beamformers are normalized according to a per-symbol power constraint, i.e. $\|\mathbf{F}\|^2 = N_{\text{T}}$ and all analog beamformers are constant modulus $|\mathbf{F}^{\text{RF}}_{i,j}|^2 = 1/N_{\text{T}}^{\text{RF}}$ for all elements (i, j) to ensure comparable results across array sizes and to reflect the per-antenna power constraint that is common in practice. In addition, the analog phase shifters have a limited resolution such that each can only apply a quantized phase shift with resolution $2^{b_{\text{phase}}}$.

We will assume a working knowledge of the beam management stages and will only briefly highlight the key equations. Readers are encouraged to review (Dreifuerst and Heath 2023b,a; Giordani et al. 2019) for an in-depth review of these processes in 5G. Given the identified focus on beam management for 6G (Lin 2023), we will assume the same three phases (SSB, CSI-RS, and feedback) as in 5G. We also assume UEs operate similarly to sub-6GHz 5G implementations, although such simplifications are only assumed to as a simplification to focus on the BS operation. Much of the beam training and hybrid operations are not specified in the standards, so we will assume a formulation similar to classical hybrid implementations found in literature i.e. (Heath et al. 2016; Singh and Joshi 2021; Wu et al. 2018).

3.3.1 Synchronization Signal Burst

The SSB process is an initial access procedure where a limited number of broadcast beams are transmitted periodically to allow UEs to connect to the network (Giordani et al. 2019; Dreifuerst and Heath 2023b). Although the first SSB codebooks were designed for wide-area coverage, the concepts of “narrow” and “wide” beams do not make sense in multipath environments, especially with extremely large arrays. Instead, the SSB and CSI-RS codebooks should be designed in a hierarchical format to find more precise

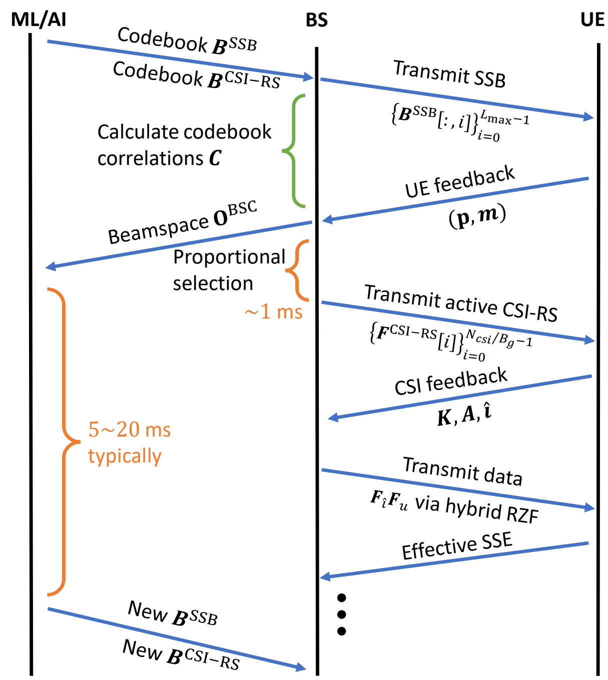

Figure 3.1: A timing diagram indicating the connection between an ML/AI codebook algorithm, the BS operation, and the UE role in a downlink configuration. Deep neural networks are capable of being deployed in near real-time roles like codebook generation with time spans on the order of 5-20ms. Real-time operations like precoder determination and feedback quantization are challenging due to the strict $\leq$ 1ms latency requirements.

beamformers that improve metrics such as the received signal power or data rate. With this in mind, we overview the SSB signal model.

The SSB signals are transmitted over a consistent band K^{SSB} and predefined timeslots T_i^{SSB} for each SSB i up to the total SSB size $L_{\max}$. The received signal during SSB reception before combining can be simplified from (3.2) to a single-stream representation of the demodulation reference signals $s_{t,k}^{\text{DMRS}} \in \mathbb{C}$ as

$$\mathbf{y}_{u,t,k}^{\text{SSB}_i} = \frac{1}{\sqrt{KN_{\text{T}}}}\mathbf{H}_{u,t,k}\mathbf{f}_i s_{t,k}^{\text{DMRS}} + \mathbf{N}_{u,t,k} \tag{3.3}$$

where $\mathbf{f}_i \in \mathbf{B}^{\text{SSB}}$ is the analog beamformer controlled from a single digital port. Note that a digital beamformer can also be used but it is most reasonably achieved through an analog rank-1 beamformer. While the limited resolution of the phase shifters b_{phase} in the hybrid array has a noticeable effect during multi-stream precoder design, the effects are almost imperceptible and nearly neglectable during rank-1 beamforming, even with low-resolution phase shifters (Alouzi et al. 2023). Because the focus is on the BS operation, we will assume the UE performs digital combining to maximize the SSB reference signal received power as

$$\mathrm{RSRP}_{i,u} = \sum_{k\in K^{\mathrm{SSB}}} \sum_{t\in T_i^{\mathrm{SSB}}} \left\| \mathbf{y}_{u,t,k}^{\mathrm{SSB}_i} \right\|^2 . \tag{3.4}$$

It is common at high frequencies to employ beam training at the UE instead of digital combining, but we leave such a scenario to future work as it is unclear whether upper-mid band 6G will operate more like sub-6GHz 5G or millimeter-wave 5G. Similarly, it is unknown if additional SSB feedback will be added in 6G, but the minimum SSB feedback corresponds to the best beam index and the RSRP of that beam. This can be aggregated at the base station as

$$\mathbf{p} = \left\{ \max_i \mathrm{RSRP}_{i,u} \right\}_{u=1}^{U} \tag{3.5}$$

$$\mathbf{m} = \left\{ \arg\max_i \mathrm{RSRP}_{i,u} \right\}_{u=1}^{U} . \tag{3.6}$$

The vectors $\mathbf{p}$ and $\mathbf{m}$ correspond to the best beam index (3.6) and the corresponding RSRP (3.5) for all U users. After the SSB feedback, the BS can transmit CSI-RS to obtain more accurate beam alignment and channel estimates in the form of PMI.

3.3.2 CSI-RS

CSI-RS is the second beamforming stage where the previous SSB feedback is used to refine the beamformers to achieve a high SNR for analog beam training and channel estimation. In the current specifications (3GPP R17), the CSI-RS codebook is large, and a subset of the candidate beamformers are selected based on the SSB feedback (Dreifuerst and Heath 2023a). We define $\mathbf{B}^{\text{sub}}$ as the active CSI-RS codebook and the selection algorithm, $\text{Sel}(\cdot)$, that determines the N_{CSI} active beamformers from $\mathbf{B}^{\text{CSI-RS}}$ based on the SSB codebook and SSBRI feedback $\mathbf{m}$. The subset selection can be formalized as

$$\mathbf{B}^{\text{sub}} = \text{Sel}(\mathbf{m}, \mathbf{B}^{\text{SSB}}, \mathbf{B}^{\text{CSI-RS}}, N_{\text{CSI}}) \tag{3.7}$$

An important challenge during CSI-RS is that the selection operates on a short timescale– on the order of 1ms or less. Because of this limitation, neural networks cannot easily be employed for selection; typical neural network execution times with modern hardware are on the order of 10us to 100ms (Ren et al. 2019; Reddi et al. 2020). Historically, the CSI-RS selection simply divides the angular space of the SSB into a proportional number of beams compared to the total allocation of CSI-RS beams, although the exact selection algorithm is left up to private implementations. Given that the learned beams do not represent simple angular spaces, we propose a similar proportional selection process using the cross-correlations of the SSB and CSI-RS codebooks. In the case of DFT beams, the proposed proportional selection produces the same results as angular division but it also extends to arbitrary, learned codebooks as well. The proportional selection algorithm begins by calculating the cross-correlations $\mathbf{C}$ between each of codebook entries, which can be performed while awaiting SSB feedback

$$\mathbf{C} = \left\langle \mathbf{B}_i^{\text{SSB}}, \mathbf{B}_j^{\text{CSI-RS}} \right\rangle \quad \forall i, j. \tag{3.8}$$

Then the number of CSI-RS beams to be allocated to each SSB reference beam is determined as a proportion of the total number of CSI-RS based on how many users reported a given SSB beam selection $\mathbf{m}$

$$\mathbf{m}^{\text{sel}} = \left\{ \left\lfloor \frac{N_{\text{CSI}}}{U} \sum \mathbf{1}_{\mathbf{m}=i} \right\rceil \right\}_{i=0}^{i=L_{\max}-1}. \tag{3.9}$$

Finally, for each of the SSB beams, the $\mathbf{m}_i^{\text{sel}}$ most correlated beams are assigned to the active subset. In other words, the selection algorithm first determines how many of the total CSI-RS beams should "correspond" to each of the SSB beams, and then the most correlated CSI-RS beams for each SSB beam are selected. This can be done based on a column-wise argsort operation, the cumulative sum operation (cumsum) and vectorized selection

$$\mathbf{A}^{\text{sorted}} = \text{argsort}(\mathbf{C}) \tag{3.10}$$

$$\mathbf{v}_i = \text{cumsum}(\mathbf{m}_{:i}^{\text{sel}}) \quad \forall i \in \{0, 1, ..., L_{\max} - 1\} \tag{3.11}$$

$$\mathbf{B}_{\mathbf{v}_{i-1}:\mathbf{v}_i}^{\text{sub}} = \mathbf{B}^{\text{CSI-RS}}[\mathbf{A}_{0:\mathbf{m}_i^{\text{sel}}}^{\text{sorted}}]. \tag{3.12}$$

The result is a subset of N_{CSI} CSI-RS beams that are proportionally selected to be highly correlated with the active SSB beams reported.

Before discussing the received CSI-RS, it is important to understand the role of logical ports when it comes to hybrid arrays and CSI-RS. Logical ports simply correspond to a set of ports that transmit a signal over a consistent channel. It does not correspond to a specific or unique subset of analog or digital ports in a hybrid system, but rather, allows for a separation between the hardware and a representation of the channel. Although this logical-physical port separation is flexible, this formulation leaves many implementation details out of the specifications. In order to follow traditional hybrid systems, while still keeping the flexibility, we associate three parameters with the CSI-RS: N_{P} is the total number of ports that can be configured which is typically the digital dimension of the hybrid array, N_{CSI} is the number of CSI-RS beamformers to be transmitted, and an additional parameter B_{g} which is the number of ports and beams assigned to one CSI-RS allocation. The connection between the parameters is that B_{g} CSI-RS beams are transmitted within one resource so that $\lceil N_{\mathrm{CSI}}/B_{\mathrm{g}} \rceil$ total CSI-RS resources must be allocated and up to $\lfloor N_{\mathrm{P}}/B_{\mathrm{g}} \rfloor$ users can be multiplexed together during the subsequent data transmission. The separation here allows for "assigning" portions of the logical array of size B_{g} to each user and users can select from $\lceil N_{\mathrm{CSI}}/B_{\mathrm{g}} \rceil$ resources to feedback beamformed channel estimates from.

With the active CSI-RS codebook determined, the CSI-RS process involves transmitting training symbols $\mathbf{s}^{\mathrm{tr}}_{t,k} \in \mathbb{C}^{B_{\mathrm{g}} \times 1}$ for channel estimation using each selected beamforming codeword, grouped into B_{g} beams i.e. $\mathbf{F}_i \subseteq \mathbf{B}^{\mathrm{sub}}$ for each CSI-RS allocation i. Note that the symbols and beamformers may be one-hot or code-division multiplexed across the time-frequency resources of the CSI-RS depending on the BS-selected settings. The received signal for each UE is then

$$\mathbf{y}^{\text{CSI-RS}_i}_{u,t,k} = \frac{1}{\sqrt{K N_{\mathrm{T}}}} \mathbf{H}_{u,t,k} \mathbf{F}_i \mathbf{s}^{\mathrm{tr}}_{t,k} + \mathbf{N}_{u,t,k}. \tag{3.13}$$

From the received signal, the UE can determine the RSRP, estimated SNR, and select the CSI-RS indicator $\hat{\mathbf{i}}$ as

$$\mathrm{RSRP}_{i,u}^{\text{CSI-RS}_i} = \max_{n_\mathrm{r}} \sum_{k \in K^{\text{CSI-RS}_i}} \sum_{t \in T^{\text{CSI-RS}_i}} \left\| \mathbf{y}_{u,t,k}^{\text{CSI-RS}_i} \right\|^2 \tag{3.14}$$

$$\mathrm{SNR}_{i,u} = \max_{n_\mathrm{r}} \frac{1}{K N_\mathrm{T}} \sum_{k \in K^{\text{CSI-RS}_i}} \sum_{t \in T^{\text{CSI-RS}_i}} \frac{\left\| \widehat{\mathbf{H}_{u,t,k}\mathbf{F}_i} \right\|^2}{\sigma^2} \tag{3.15}$$

$$\mathbf{F}_{\hat{i}_u} = \arg\max_i \mathrm{SNR}_{i,u} \tag{3.16}$$

The noise power, $\sigma^2 = \mathbb{E}[\|\mathbf{N}_{u,t,k}\|^2]$ can be estimated from the system parameters and zero-power CSI-RS. Compared to the SSB model (3.3), the CSI-RS received model is multi-layer due to the need for sounding each of the B_g ports. The CSI-RS therefore must have at least N_R beams for a full-rank beamformed channel, assuming $N_\mathrm{R} \leq N_\mathrm{T}$. There is a significant benefit to using groupings of a small number of beams to sound the channel as the overhead for both CSI-RS transmission and feedback quantization are reduced, and users can be multiplexed using different analog beams and well-designed digital precoders.

The SNR (3.15) impacts the channel estimation efficacy and is useful for determining the initial modulation and coding scheme (MCS) for data transmission. The beamformed channel estimate is defined for a channel estimation algorithm $w(\cdot)$ as

$$\widehat{\mathbf{H}_u\mathbf{F}_{b_u,\hat{i}_{b_u,u}}} = w(\mathbf{y}_u^{\text{CSI-RS}_{\hat{\mathbf{i}}_u}}, \mathbf{s}^{\mathrm{tr}}). \tag{3.17}$$

There is no standardized channel estimation algorithm, but, given the limited time and power budget of mobile UEs, a least squares method is assumed herein. The resulting channel estimate at the UE is frequency selective over a set of frequency subbands $b \in \{1, ...\mathrm{S_B}\}$. In the final step, the channel estimates are quantized according to a feedback codebook and transmitted to the BS (Dreifuerst and Heath 2023b; Qin and Yin 2023).

3.3.3 Feedback

CSI feedback is a highly configurable process to achieve a careful balance of overhead and CSI accuracy. In 5G there were two categories of feedback: type-I and type-II (Dreifuerst and Heath 2023b; Qin and Yin 2023). In 3GPP Release 16 an additional format of type-II called "enhanced type-II" was standardized that included additional, frequency-domain compression. In general, the feedback formats are methods for quantizing the channel estimate according to a shared feedback codebook $\mathbf{B}^{\mathrm{FB}}$. Type-II in particular is designed for high-resolution, MU-MIMO settings where the channel information is quantized as a set of L_{CSI} oversampled DFT components, with oversampling factors $O_{\mathrm{H}} \times O_{\mathrm{V}}$. We refer readers to (Qin and Yin 2023) for a general description of the feedback formats and (Dreifuerst and Heath 2023c) for a fast implementation of type-II quantization that is important for large arrays and high resolution feedback considered in X-MIMO.

The BS receives a set of codebook indices $\mathbf{A} \in \mathbb{W}^{\mathrm{S_B} \times N_{\mathrm{R}} \times L_{\mathrm{CSI}}-1}$ and cophasing factors $\mathbf{K} \in \mathbb{C}^{\mathrm{S_B} \times N_{\mathrm{R}} \times L_{\mathrm{CSI}}-1}$ that enable a reconstruction of the beamformed channel. Channels are reconstructed in the same way as (Dreifuerst and Heath 2023a) i.e.

$$\widehat{\mathbf{H}_{u,b,n_{\mathrm{r}}}\mathbf{F}_{b_u,\hat{i}_{b_u,u}}} = \sum_{\ell=1}^{L_{\mathrm{CSI}}} \mathbf{K}_{b,n_{\mathrm{r}},\ell} \mathbf{B}^{\mathrm{FB}}_{\mathbf{A}_{b,n_{\mathrm{r}},\ell}} \quad \forall b,\ n_{\mathrm{r}}. \tag{3.18}$$

The quantization error, therefore, can be reduced with large L_{CSI}, but this introduces additional overhead. In addition, the frequency domain quantization introduces overhead that increases with a larger number of subband elements $\mathrm{S_B}$. The actual overhead is compressed so that the frequency-domain scaling is minimal, although the overhead of this signaling is a significant part of the PMI section of the feedback report. In addition, the feedback will include information describing the rank indicator (RI) R, the best CSI-RS beam index $\hat{i}_u$ (CRI), and other subfields not critical to this work. The base station should then use the information from both the beam training and channel feedback to design precoders to best serve the users.

3.3.4 Data transmission

In the final steps, the BS uses the CSI acquired through the UE feedback to determine the precoders to serve a set of users. The exact process is not standardized; we take a typical approach here inspired by MIMO fundamentals (Heath Jr. and Lozano 2018). The analog beamforming for the hybrid array is controlled through the CRI from the feedback packet while the digital precoder is used to mitigate interference between users and streams. A typical precoder formulation is an RZF precoder determined to maximize the signal-to-leakage noise ratio (Heath Jr. and Lozano 2018, section 9.9). The RZF precoder with perfect channel state information is

$$\mathbf{F}_{u,t,k} = \frac{(\sum_{i=0}^{U-1} \mathbf{H}_{i,t,k}^{*}\mathbf{H}_{i,t,k} + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^{2}])^{-1}\mathbf{H}_{u,t,k}^{*}}{\left\|\mathbf{H}_{i,t,k}^{*}\mathbf{H}_{i,t,k} + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^{2}])^{-1}\mathbf{H}_{u,t,k}^{*}\right\|}. \tag{3.19}$$

In the case of realistic, hybrid arrays the perfect channel state information $\mathbf{H}$ is replaced with the selected beamformed channels, redefined for ease of notation as $\widehat{(\mathbf{HF})} \equiv \{\widehat{\mathbf{H}_{u,b,n_{\mathrm{r}}}\mathbf{F}_{\hat{i}_u}}\}_{u,b,n_{\mathrm{r}}} \in \mathbb{C}^{U\times \mathrm{S_B}\times N_{\mathrm{R}}\times B_{\mathrm{g}}}$. The beamformed channels arise from the beam grouping and port assignments from Section 3.3.2. Without loss of generality, we assume the first $U_a \leq N_{\mathrm{Td}}/B_{\mathrm{g}}$ users are allocated for data transmission so that the analog $\{\mathbf{F}_{\hat{i}_u}\}$ and digital precoders $\{\mathbf{F}_{u,t,k}\}$ are designed in a subarray-formulation (Heath et al. 2016) as

$$\mathbf{F}^{\mathrm{RF}} = [\mathbf{F}_{\hat{i}_0}, \mathbf{F}_{\hat{i}_1}, ..., \mathbf{F}_{\hat{i}_{U_a}}] \tag{3.20}$$

$$\mathbf{F}_{t,k} = \begin{bmatrix} \mathbf{F}_{0,t,k} & 0 & ... & 0 \\ 0 & \mathbf{F}_{1,t,k} & 0 & ... \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & ... & \mathbf{F}_{U_a,t,k} \end{bmatrix}. \tag{3.21}$$

This formulation neatly unifies hybrid arrays, traditional precoding, and beam management techniques with the received signal model from (3.2).

On the receiver side, the UEs also perform a combining strategy based on the embedded DMRS within the signal allowing for decoding each data stream without knowing the exact precoder. In this work, we will assume an LMMSE receiver is used to maximize the signal-to-interference noise ratio for any precoder, although the actual SINR and subsequent data rate calculations are not actually performed by the UE but are necessary to evaluate the achievable spectral efficiency in simulation. The choice of LMMSE receiver here makes sense to ensure allocated users have high performance for receiving the signal while still balancing the computational cost. The SINR at resource element (t, k) for a given user u and data stream r with equal power allocation per user is

$$\mathbf{I}_{u,t,k} = \left(\sum_{i=0}^{U-1} \mathbf{H}_{u,t,k} \mathbf{F}_{b_u, \hat{i}_{b_u,u}} \mathbf{F}_{i,t,k} (\mathbf{H}_{u,t,k} \mathbf{F}_{b_u, \hat{i}_{b_u,u}} \mathbf{F}_{i,t,k})^* + U N_{\mathrm{T}} \mathbb{E}[\mathbf{N}_{u,t,k}^2] \right)^{-1} \tag{3.22}$$

$$\mathrm{SINR}_{u,t,k,r} = \frac{1}{K} \frac{(\mathbf{H}_{u,t,k} \mathbf{F}_{b_u, \hat{i}_{b_u,u}} \mathbf{F}_{u,:,r})^* \mathbf{I}_{u,t,k} \mathbf{H}_{u,t,k} \mathbf{F}_{b_u, \hat{i}_{b_u,u}} \mathbf{F}_{u,:,r}}{1 - (\mathbf{H}_{u,t,k} \mathbf{F}_{b_u, \hat{i}_{b_u,u}} \mathbf{F}_{u,:,r})^* \mathbf{I}_{u,t,k} \mathbf{H}_{u,t,k} \mathbf{F}_{b_u, \hat{i}_{b_u,u}} \mathbf{F}_{u,:,r}}. \tag{3.23}$$

The SINR expression (3.23) is a ratio between the signal power fo the desired data stream (u, r) compared to the noise and interference power of all other streams in (3.22).

We evaluate the performance in the network by the sum spectral efficiency, which is reasonable when there is not more complicated scheduling, fairness constraints, and resource allocation. The sum spectral efficiency can be multiplied by the bandwidth to get the sum rate. Assuming Gaussian signaling and treating interference as Gaussian noise, the achievable spectral efficiency, $\mathrm{SE}_{u,t}$, is

$$\mathrm{SE}_{u,t} = \sum_{k=0}^{K-1} \sum_{r=0}^{R-1} \log_2(1 + \mathrm{SINR}_{u,t,k,r}). \tag{3.24}$$

The SSE is a sum over the active users, but it does not account for the overhead of beam management or CSI acquisition. In the final step, we consider the ESSE, which accounts for the overhead due to beam training by removing the corresponding time/frequency resources due to training and feedback of the entire beam management system ($T_{\mathrm{BM}}, K_{\mathrm{BM}}$) from the spectral efficiency calculation

$$\mathrm{ESSE} = \sum_{u=0}^{U-1} \sum_{\substack{t \notin T_{\mathrm{BM}} \\ k \notin K_{\mathrm{BM}}}} \sum_{r=0}^{R-1} \log_2(1 + \mathrm{SINR}_{u,t,k,r}). \tag{3.25}$$

Although presented in a linear fashion, the SSB, CSI-RS, feedback, and data transmission typically occur over separate resources so that there are no periods of waiting in a full-buffer scenario. With the system model defined, we now present the core machine learning architecture, X-BM that integrates ML within the codebook design process. The proposed strategy integrates in a realistic fashion such that the system does not require additional side information and the overall architecture is designed to meet the timing constraints implicit within beam management.

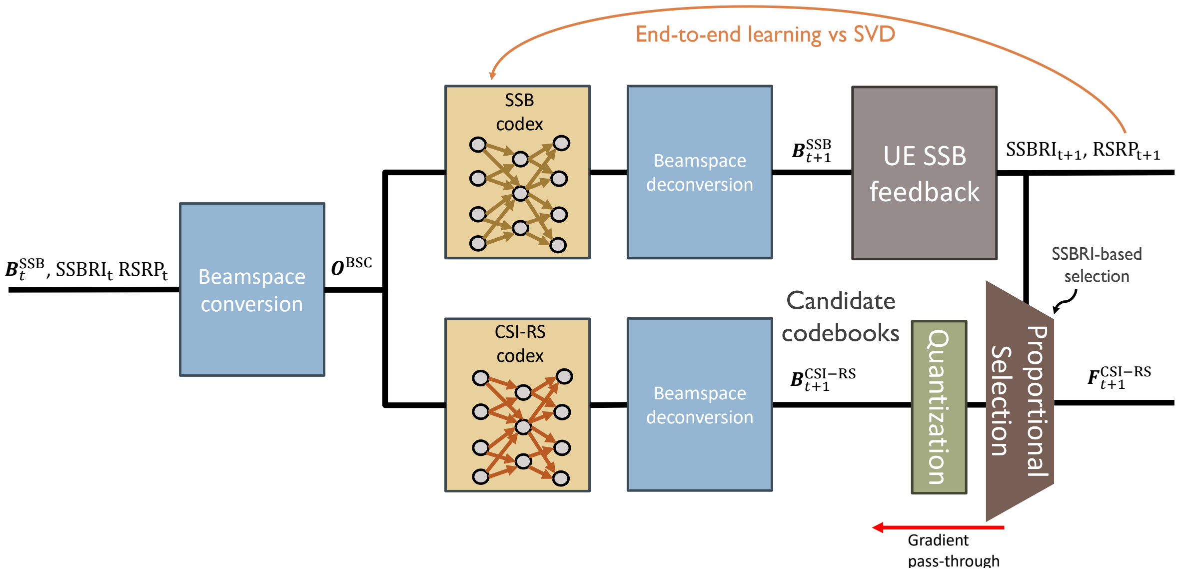

Figure 3.2: A depiction of the X-BM algorithm and gradient propagation. The beamspace is prepared using the previous SSB codebook and feedback and fed into two parallel networks that generate the new SSB codebook and CSI-RS codebook. The CSI-RS codebook is then quantized and a proportional selection according to the new SSB feedback is used to determine the active CSI-RS codebook. Gradients are calculated according to the difference between the UE-received power and the achievable SVD received power.

3.4 Neural codebook design

We propose a novel approach for jointly designing the SSB and CSI-RS codebooks using AI/ML. The primary challenges with integrating machine learning into wireless domains ultimately come down to the tight timing constraints (see Figure 3.1) and generalizing across varying numbers of users and system configurations. Timing constraints are not particularly concerning for SSB codebook generation, as the periodicity is typically 20ms with a minimum of 5ms. In comparison, the CSI-RS codebook needs to be available within 1ms following the SSB feedback, however, the CSI-RS selection depends on the SSB feedback. It is extremely challenging for modern hardware to employ deep neural networks with such tight, sequential timing constraints. Therefore, we propose the neural architecture, shown in Figure 3.2, to output the SSB codebook and a set of candidate CSI-RS codebooks in parallel at each SSB period. Then we employ the proportional selection technique from (3.8)-(3.12) to determine the CSI-RS codebook from the candidate options, allowing the joint algorithm to overcome the timing limitation and still incorporate the SSB feedback. In preparation for 6G, this operation is similar to current 5G CSI-RS codebooks but extends to arbitrary and time-varying SSB and CSI-RS codebooks. The entire structure is trained in an end-to-end fashion with the goal of minimizing the power loss with respect to singular value decomposition beamforming, which is the power-maximizing strategy for a single user with perfect CSI and rank-1 beamforming. This ensures that we do not have to define a specific beamforming pattern train against, but rather have the network learn patterns that are close to power-maximizing.

The input to X-BM is prepared by converting the prior SSB codebook and SSB feedback to the beamspace. The input only contains previously known users, so the learning algorithm is partially regularized by new users who join the network in the upcoming SSB cycle and therefore are not included in the prior observation. The beamspace is a simple and highly extensible conversion where beamformers are converted to the angular domain with fixed resolution. This format is beneficial for translating beamformer information over time, antenna geometries, and UE numbers because the beamspace is largely unaffected by these parameters (Sayeed 2002; Wen et al. 2018; Dreifuerst and Heath 2023a). This ensures that the input from a set of previous beamformers represents a consistent two-dimensional grid of projections. This step was originally developed in our prior work (Dreifuerst and Heath 2023a).

First, we define an angular transformation matrix for antenna size N_{s} and N_{m} angular directions as a series of array responses

$$\boldsymbol{\theta}_{N_{\mathrm{m}}} = \frac{1}{\pi}[0, 1, ..., N_{\mathrm{m}} - 1]^T \tag{3.26}$$

$$\mathbf{U}_{N_{\mathrm{s}},N_{\mathrm{m}}} \triangleq [\mathbf{a}_{N_{\mathrm{s}}}(\boldsymbol{\theta}_0), ...\mathbf{a}_{N_{\mathrm{s}}}(\boldsymbol{\theta}_{N_{\mathrm{m}}-1})]. \tag{3.27}$$

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

$$\mathbf{F}_i^{\mathrm{SSB}} = [\mathbf{B}^{\mathrm{SSB}}_{0:N_{\mathrm{X}}^{\mathrm{RF}},i}, ..., \mathbf{B}^{\mathrm{SSB}}_{(N_{\mathrm{X}}^{\mathrm{RF}}-1)N_{\mathrm{Y}}^{\mathrm{RF}}:N_{\mathrm{X}}^{\mathrm{RF}}N_{\mathrm{Y}}^{\mathrm{RF}},i}] \tag{3.28}$$

$$\mathbf{O}_i^{\mathrm{BSC}} = \mathbf{U}^*_{N_{\mathrm{X}}^{\mathrm{RF}},N_{x0}}\mathbf{F}_i^{\mathrm{SSB}}\mathbf{U}_{N_{\mathrm{Y}}^{\mathrm{RF}},N_{y0}} \quad \forall i. \tag{3.29}$$

The codebook must be reshaped from a vector of size $N_{\mathrm{T}}^{\mathrm{RF}}$ to the planar dimensions $N_{\mathrm{X}}^{\mathrm{RF}} \times N_{\mathrm{Y}}^{\mathrm{RF}}$ in (3.28) before the beamspace conversion in (3.29) that produces the beamspace observation, $\mathbf{O}^{\mathrm{BSC}}$. Note that the beamspace conversion is a reversible operation via the pseudo inverse so long as the number of samples in each direction is equal to or greater than the number of antennas, so we also train the network to predict the beamspace of the desired output, rather than direct beamforming coefficients to aid the learning process. The predicted beamformers are obtained within the training loop, 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

$$\mathbf{O}^{\mathrm{BSC}} = \left\{\left[\mathbf{O}^{\mathrm{BSC}}[i], \sum \mathbf{1}_{\mathbf{m}=i}, \mathbf{1}^T_{\mathbf{m}=i}\mathbf{p}\right]\right\}_{i=0}^{L_{\max}-1}. \tag{3.30}$$

This input is generated during each SSB period and fed into the X-BM algorithm as shown in Figure 3.2.

The overall system is designed and trained in an end-to-end fashion with respect to maximizing the RSRP. End-to-end training changes the learning goal with data-driven methods to maximize or minimize a system-level metric, rather than training toward a specific output (Felix et al. 2018). In particular, we do not identify the exact beamformers we want the neural network to learn, instead relying on the network to learn the beamformers that achieve the highest RSRP. It is difficult to exactly define an "optimal" beamformer formulation in general for the SSB and CSI-RS steps. This arises from the fact that even with perfect CSI, the broadcast beamformers that achieve the highest SNR with a limited number of beams depend heavily on the system parameters, user channels, and number of beamformers used. By employing end-to-end learning, we instead train the neural network to learn beamformers that achieve the highest performance relative to single-user singular value decomposition beamforming with perfect CSI. This formulation is even more advantageous in the initial access scenario where not all users are known by the system prior to codebook generation. The general learning framework is shown in Figure 3.2 with a loss function comparing the SVD power and the reported RSRP for each user as

$$\mathbf{U}_u, \mathbf{S}_u, \mathbf{V}_u = \mathbf{H}_u \quad \forall u \in \{1, ...U\} \tag{3.31}$$

$$\mathcal{L}(\mathbf{S}, \mathbf{p}) = \frac{1}{U} \sum_u (10 \log_{10}(\mathbf{S}^2_{u,0}) - 10 \log_{10}(\mathbf{p}^2_u)) \tag{3.32}$$

where $\mathbf{p}$ is the RSRP from (3.5) or (3.14) for the SSB or CSI-RS codebooks. Gradients are calculated from the loss function (3.32) throughout the network and parameters are updated using an Adam optimizer (Kingma and Ba 2014) with cyclic learning rate

Table 3.1: X-BM neural architectures

Layer	**Primary Parameter**	**Act.**	**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	0.1 Rate		$(\lceil(N_{x0}+2)/4\rceil)(\lceil(N_{y0}+2)/4\rceil)(320)$
Fully Conn.	$2N_{x0}N_{y0}N$ Neurons		$2N_{x0}N_{y0}N$
Reshape			$N_{x0} \times N_{y0} \times 2N$

scheduling between 10^{-3} and 10^{-5}.

One challenge of codebook design for X-MIMO is the massive physical array size, which results in a significant number of parameters for generating codebooks. We mitigate this with efficient neural structures like convolutional layers (See Table 3.1) that share parameters and efficiently operate over the beamspace representation (3.29). In fact, this formulation allows the neural network to learn from a universal pattern that allows for training and testing on any array geometries, so long as the beamspace dimensions are equal or greater to the antenna dimensions to prevent aliasing (Dreifuerst and Heath 2023a). We examine how the network generalizes across geometries in Figure 3.6. It is noteworthy that further parameter reduction can be achieved using convolution transpose layers instead of a fully connected layer within the network, although we found this caused a slight performance drop in our previous work (Dreifuerst and Heath 2023a).

3.5 Data generation framework

The system-level simulation requires two stages: 1) spatially and consistent channel generation for multiple users and 2) beam management integration of the iterative steps for initial access, beam refinement, feedback, and data transmission. We start by generating raytraced channels using Sionna (Hoydis et al. 2022). In our simulations, we consider two environments (labeled A and B) with 10000 potential users scattered over a BS sector site. The BS is equipped with an $N_{\mathrm{X}}^{\mathrm{RF}} = 16$, $N_{\mathrm{Y}}^{\mathrm{RF}} = 16$ planar array using a 3GPP 3D antenna model with half-wavelength spacing and $N_{\mathrm{X}} = 8$ and $N_{\mathrm{Y}} = 4$ with $b_{\mathrm{phase}} = 2$ bit phase shifters facing the serving sector region. Each user has $N_{\mathrm{R}} = 4$ antennas tuned for a 10GHz carrier frequency in environment A and tuned for a 20GHz carrier frequency in environment B to consider low-band and high-band frequencies for upper-mid bands (Holma et al. 2021). Users are assumed to have random mobility patterns with an average speed of 3m/s in environment A and 10m/s in environment B and arrays oriented vertically. Environment A models the area around the Frauenkirche in Munich with a BS placed on a building 40m tall and Environment B models the area around the Arc de Triomphe in Paris with the BS at a height of 27m (Hoydis et al. 2022). The user channels are sampled at 1ms intervals over 100MHz with 270 resource blocks assuming 30kHz subcarrier spacing. After a database of channels is generated, the neural network is integrated into a processing pipeline following the system model. The neural network is trained on $40,000$ training samples with a random number of users and varied locations throughout the scene. The validation dataset is drawn from the same distribution and used to determine when the network training is complete. All test results come from newly generated channels according to the test environment, which may be the same or a different environment from the training set.

The data processing begins by first selecting a random number of active users $U \sim \mathcal{U}[4, 16]$ for the subsequent timestep and drawing that many users' channels randomly from the channel database. Then, the beamspace is calculated using the prior SSB codebook (initially set as random DFT codebook beams) with 80% of the active users reporting. The remaining 20% are assumed to be new users that are not known prior to the current SSB cycle and helps prevent the neural network from overfitting to known users. The new SSB and CSI-RS codebooks are generated and the current iteration begins with the SSB transmission and feedback ($\mathbf{p}$, $\mathbf{m}$) according to Section 3.3.1. Then, the active CSI-RS codebook $\mathbf{B}^{\text{sub}}$ is selected according to the proportional correlation selection, beams are grouped into sets of length P_{CSI} to form beamformed channels and transmitted. During data transmission, we assume the network calculates the RZF precoders based on the feedback and estimates the expected SSE for each combination of UEs, and selects the highest SSE user set. This information is used as a simple, realistic scheduling algorithm using the CSI available. The resulting SSB RSRP, CSI-RS SNR, and ESSE metrics are used to evaluate the proposed X-BM from a modular and system-level context.

3.6 Simulation results

Given the capabilities for overfitting with neural networks, we carefully evaluate the proposed system across different scenarios, array geometries, and out-of-distribution settings. From these evaluations, we seek to better understand the advantages and limitations of a neural codebook algorithm. Unless specified, the default parameters for the results are $L_{\max} = 16$ SSB beams, $N_{\text{CSI}} = 16$ CSI-RS beams, $B_{\text{g}} = 4$ logical ports or CSI-RS beams per CSI-RS resource, $N_{\text{RB}} = 24$ resource blocks occupied by a CSI-RS resource, and $\text{S}_{\text{B}} = 8$ subbands for PMI.

3.6.1 SSB Performance

The first evaluation focuses on the RSRP performance of the proposed algorithm. The RSRP is evaluated using both the SSB and CSI-RS codebooks, although the CSI-RS relies on the SSB feedback and should in general outperform the SSB beamforming given the additional information available to the system. In Figure 3.3 we show the empirical probability density function (PDF) of the reported RSRP using the proposed X-BM as well as DFT codebooks. To compare the results, we show the relative power-loss of the RSRP achieved with the codebook method compared to SVD beamforming with perfect CSI (CSI-SVD). In the first results, we show the RSRP performance with $L_{\max} = 16$, which is more than supported in sub-6GHz 5G, but less than the $L_{\max} = 64$ supported by mmWave 5G. It can be seen that the performance of X-BM is almost always within 3dB of SVD beamforming and achieves a gain of over 6dB on average compared to DFT codebooks. Especially interesting, we see that there is very little gain from the CSI-RS with $N_{\text{CSI}} = 32$ for the majority of users, with only the bottom 20% of users seeing noticeable improvements.

In Fig. 3.4, the performance with respect to different SSB codebook sizes $L_{\max}$ is plotted. An interesting result of these simulations (Figures 3.3-3.4) is that the X-BM SSB performance often outperforms even the CSI-RS performance of DFT codebooks. From a beam training perspective, there appears limited need for using CSI-RS beams when employing X-BM codebooks, because the SSB codebooks are already well-performing. Although a hierarchical search provides significant gains with DFT codebooks, it can be seen that the X-BM SSB codebook typically achieve a sufficient beam alignment. In particular, X-BM achieves better beam training RSRP with $L_{\max} = 8$ beamformers than the CSI-RS beams achieve from a 4 times larger DFT codebook, although these results

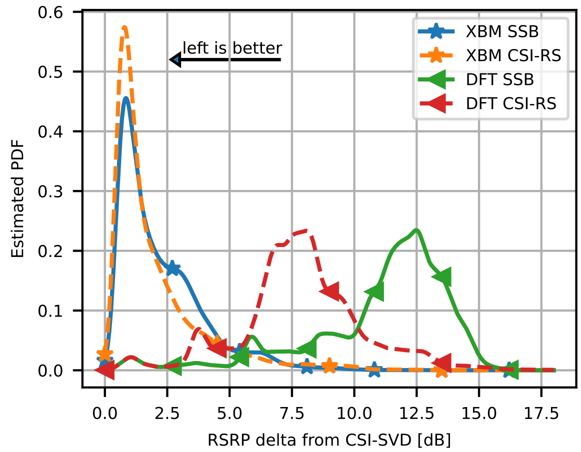

Figure 3.3: Comparison of the beamforming codebook performance relative to perfect CSI beamforming with $L_{\text{max}} = 16$ and $N_{\text{CSI}} = 16$ beams. 5000 sets of active users are drawn and each codebook method is used to select the SSB (solid lines) and CSI-RS (dashed lines) codebooks. The RSRP is then calculated and the difference from SVD beamforming is shown as an estimated density function. The X-BM codebooks are extremely effective, with more than 75% of users within 3dB of perfect CSI. Even compared to much larger DFT codebooks for CSI-RS, there is a substantial performance improvement with the proposed X-BM method.

can still be improved with a larger X-BM codebook as shown in Figure 3.4.

3.6.2 CSI-RS SNR

Following the SSB transmission, the candidate CSI-RS codebook is selected based on the SSBRI and transmitted. The first goal is to achieve a high channel estimation SNR, while also ensuring the beamformer can be employed for hybrid beamforming with digital precoding. An inherent challenge of the design and operation of CSI-RS codebooks is the dependence on the SSB codebook and the extensive overhead associated with larger active codebooks. It is often advantageous to employ a larger number of SSB beams, which require relatively low overhead, in combination with a limited number of CSI-RS

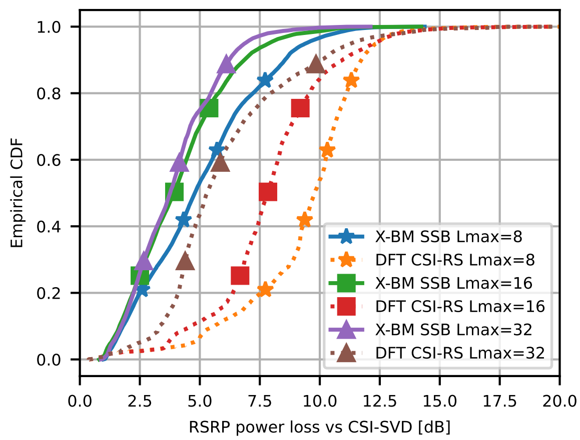

Figure 3.4: Comparison of the X-BM SSB RSRP with different sizes of codebooks ($L_{\max}$) versus CSI-RS DFT codebooks. The DFT CSI-RS codebooks always include $N_{\text{CSI}} = 32$ active beams selected based on the proportional selection algorithm with the reported SSBRI. The proposed X-BM SSB codebooks outperform 4 times larger SSB+CSI-RS DFT codebooks.

beams that consume significant wireless resources. At the same time, the SSB periodicity is higher than CSI-RS–typically about 4 times faster–so there is additional overhead to consider. Furthermore, the design challenge is not simply to maximize SNR, but also to obtain an analog precoder that performs well for the hybrid data transmission.

We begin by addressing the first question of achieving a high-gain CSI-RS codebook. In Figure 3.5 we compare the CSI-RS SNR obtained from classic DFT codebooks with varying sizes compared to the number of active users. We can see that when the number of users is much larger than the number of beams, the X-BM method performs significantly better than DFT codebooks. We can see that with just 1 beam per active user the X-BM codebooks reach a consistent performance gain of about 4dB over DFT codebooks of any size. This can be attributed to the multipath propagation environment that does not match DFT beamformers in general.

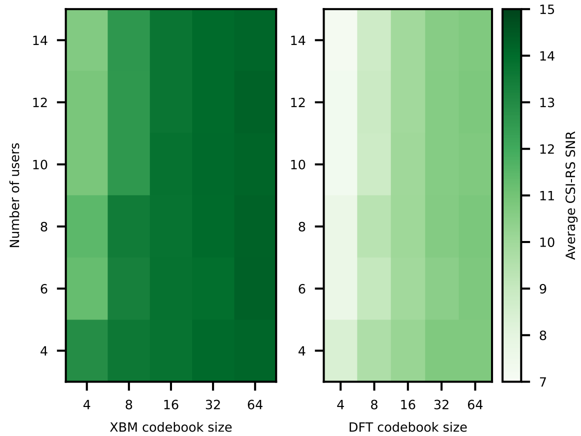

Figure 3.5: A heatmap of the average SNR during CSI-RS reception with different numbers of active CSI-RS beams N_{CSI} and active users U. While increasing the DFT codebook size provides a significant gain for more than 4 users, the rich multipath environment enables better performance with X-BM codebooks, even with 1/16 as many CSI-RS beams. This is a limitation of DFT beamformers in low- and mid-band environments.

3.6.3 Geometry translation

One of the beneficial aspects of the X-BM structure is a natural translation across array geometries. In particular, the neural network size and parameters do not depend on the antenna array, so the same neural network can be employed or trained across geometries. Figure 3.6 shows a comparison of the performance of geometry translation, assuming the arrays are placed in the same scene and location. We focus on the more difficult setting of training a neural network on data from a smaller array than it is tested on. It can be seen that there is a drop in performance when testing models trained on 8×8 arrays in 16×16 settings, although the performance is still better than the DFT codebooks in Figure 3.3. In general, this is a limitation of deep learning in that does not automatically generalize well, even with carefully designed formulations. While the performance loss is noticeable, the benefit to the X-BM architecture is that the same neural network can

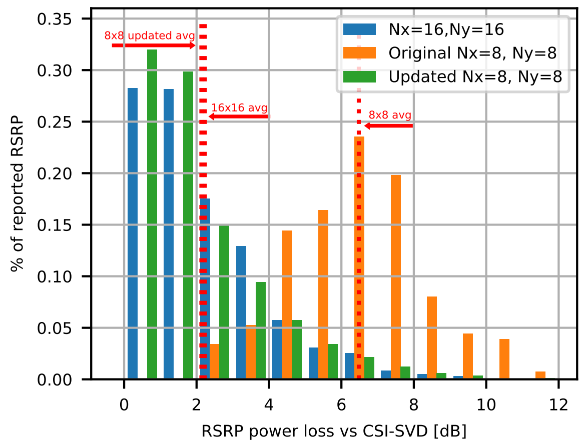

Figure 3.6: A comparison of two X-BM networks for geometry translation. The X-BM models are trained using data from arrays equipped with $N_\mathrm{X} = \{8, 16\}$ and $N_\mathrm{Y} = \{8, 16\}$ and beamspaces of size $N_{x0} = 16, N_{y0} = 16$, but both are evaluated with $N_\mathrm{X} = 16$, $N_\mathrm{Y} = 16$ and $L_\mathrm{max} = 16$. Initially, there is a performance drop of about 4dB, but fine-tuning restores the performance to that of a model trained specifically for the array geometry. This highlights the effectiveness and generalization capability of the beamspace conversion.

be trained with data from different array sizes. This allows network operators to only need to collect data once and then fine-tune after selecting the hardware to maximize performance. To show the fine-tuning performance, we include the original 8×8 model that has been retrained on the 16×16 data for 30s, highlighting how effective fine-tuning is for resolving the geometry translation loss. In this setting, fine-tuning is performed by simply training on new data with the same training steps, however, we will consider fine-tuning on imperfect CSI in Section 3.6.5.

3.6.4 Feedback

The previous results have shown that ML codebooks are extremely effective for beam training, but beam alignment is not the ultimate goal of wireless networks. Effective beam training naturally ensures that a strong rank-1 wireless link exists, however, the actual MU-MIMO performance is less clear from the beam alignment results. In this subsection, we incorporate type-II feedback and multi-stream data transmission to determine the system-level impact of ML codebooks. We first address the general impact of the CSI-RS parameters for hybrid arrays, then look at a comparison between traditional codebooks and X-BM codebooks. Throughout these results, the most important metric is always the effective sum spectral efficiency, which describes the actual performance after accounting for the overhead.

First, we investigate the feedback parameters using traditional DFT codebooks and hybrid arrays in Fig. 3.7. These results show an interesting result: using higher rank CSI-RS requires more feedback and typically reduces the performance compared to lower rank CSI-RS. This means there is little benefit to configuring the system with many CSI-RS ports. In fact, MU-MIMO with at least one user receiving 4 data layers is unlikely to maximize performance so even smaller beam groupings can be used. These results are specific to the simulated scenario, which includes many UEs to multiplex between and encourages scheduling more users with lower rank channels and less overhead. The total overhead depends on the number of CSI-RS resources allocated $N_{\mathrm{CSI}}/B_{\mathrm{g}}$, the beams per group B_{g}, and the feedback and overhead associated with each CSI-RS.

In the next investigation, we compare how the SSB and CSI-RS codebooks impact the overall system ESSE. The results in Figure 3.8 are compared again using DFT and X-BM codebooks with three feedback settings from the previous results, ($L_{\mathrm{CSI}} = 2, B_{\mathrm{g}} = 2$), ($L_{\mathrm{CSI}} = 2, B_{\mathrm{g}} = 4$), and ($L_{\mathrm{CSI}} = 4, B_{\mathrm{g}} = 4$). It can be seen that the X-BM codebooks result in roughly a 10% effective spectral efficiency improvement over DFT codebooks using the standard feedback formats. Furthermore, the gain from using limited feedback is increased with X-BM codebooks due to the more effective beam alignment.

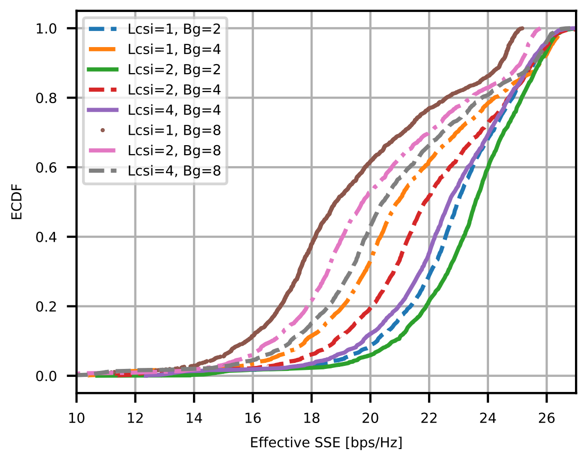

Figure 3.7: An empirical CDF chart of the ESSE performance with PMI type-II using different ratios of feedback resolution L_{CSI} and CSI rank B_{g}. Each ratio is plotted with different line formats. It can be seen that there is no gain from using larger beam groupings which require more overhead, more feedback, and reduces the number of UEs that can be multiplexed.

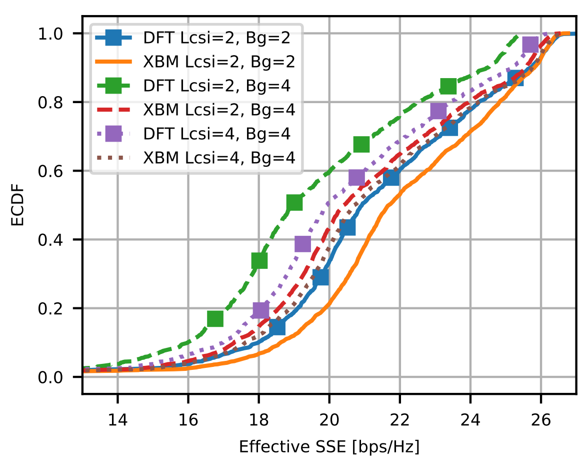

Figure 3.8: ESSE performance with CSI parameters ($L_{\text{CSI}} = 2, B_{\text{g}} = 2$), ($L_{\text{CSI}} = 2, B_{\text{g}} = 4$), and ($L_{\text{CSI}} = 4, B_{\text{g}} = 4$). The X-BM codebooks provide a noticeable improvement over DFT codebooks, especially for the lowest 20% of users. While the significant gain in beam training does result in improved ESSE, it is not expected to see gains on the order of 4 times larger ESSE, especially with greedy scheduling.

3.6.5 Model updating

For the final evaluation, we evaluate how the model performs as a result of distribution shift. We consider the case that the model is initially trained in the first environment, but then the scene changes significantly causing the channel distributions to be out-of-distribution (Ben-David et al. 2010) or arising from a new context. We evaluate the model in the new environment (environment B) and then again after retraining the model using only the available CSI at the BS. In particular, we retrain on potentially erroneous CSI to characterize how effective live-model updating is using imperfect CSI. Because of the lack of perfect CSI, we employ an unsupervised learning strategy to maximize the RSRP (or minimize the negative RSRP)

$$\mathcal{L}(\mathbf{p}) = -\frac{1}{U}\sum_{u} 10\log_{10}(\mathbf{p}_u)^2. \tag{3.33}$$

This formulation is less stable as the training only seeks to maximize the RSRP without a reference point so a low RSRP is more penalizing, even if that UE cannot achieve a much higher RSRP due to the location/channel. Additionally, the gradients are calculated with respect to the quantized, estimated beamformed CSI which can include estimation and quantization error. At the same time, the training considers a very practical scenario where only the available information obtained from a live network is used for fine-tuning.

The results of the fine-tuning are shown in Figure 3.9 which characterizes the difference between a site-specific model, the site-agnostic model, and the fine-tuned model performance between the scenarios. It can be seen that there is a significant drop in performance when deploying the initial model (site-agnostic) in the new environment, but that using the imperfect CSI and retraining with 20000 samples for 30s can provide an improvement, especially for the bottom 50% of users. The imperfect CSI and unsu-

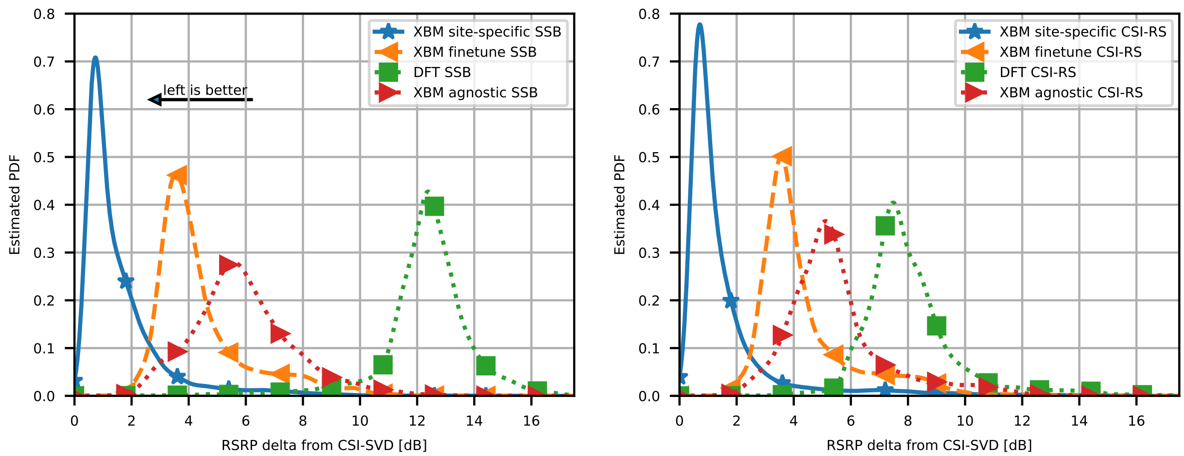

(a) SSB

(b) CSI-RS

Figure 3.9: Empirical PDF comparison of the SSB RSRP (a) and CSI-RS RSRP (b) for the X-BM model deployed in environment B and training on data from environment A at a different carrier frequency and user distribution. The site-specific model (−), fine-tuned (−−), and site-agnostic model (··) correspond to a model trained directly for environment B, a model trained on environment A and retrained on environment B, and the original model from environment A without retraining. Even trained on a different environment, the proposed X-BM method outperforms DFT codebooks and transfers to new environments efficiently.

pervised fine-tuning strategy, however, is not able to recover the results that could be achieved with perfect CSI shown by the site-specific results. The flexibility of deploying pre-trained models and then quickly fine-tuning is a valuable benefit for network operators, even during cyclic changes in distributions i.e. rush hour or during sporting events.

3.7 Conclusion

In this chapter, we introduced an extended AI/ML algorithm that designs SSB and CSI-RS codebooks for X-MIMO systems. The strategy builds upon our beamspace learning algorithm through a parallel-codebook strategy that integrates the timing and computational constraints of beam management. We use an end-to-end training strategy to back-propagate gradients through the hierarchical beam search to learn effective codebooks. The X-BM algorithm is shown to perform exceptionally well in raytraced environments with a typical beam alignment gain of 6dB over traditional methods. Even reducing the X-BM codebooks to a single-stage search during SSB still achieves a better beam alignment than two-tier DFT codebook search. At the same time, the ESSE is improved by less than 10% on average. This shows that training the neural network to maximize beam training power results in strong rank-1 beamformers that are not necessarily advantageous for MU-MIMO communications.

Chapter 4

Neural Codebook Design for Network Beam Management

4.1 Overview

In this chapter[1], we finalize the beamspace learning paradigm to mitigate multi-cell interference and maximize the achievable rate. Using this formulation, the learning focuses on the network-level ESSE metric instead of beam training power which ultimately results in multi-rank MU-MIMO channels capable of supporting higher degrees of spatial multiplexing. This chapter in particular showcases how end-to-end learning can improve realistic, multi-cell beam management by a significant degree without requiring massive feedback or modifying UE operation.

[1] This chapter is based upon my work in (Dreifuerst and Heath 2024b): Dreifuerst, R. M. and Heath, R. W. (2024b). Neural codebook design for network beam management. *Submitted to IEEE Trans. Wireless Commun.*.

4.2 Introduction

Multiple-input multiple-output (MIMO) systems have played a critical role in supporting high-speed wireless communications (Heath Jr. and Lozano 2018). Between spatial multiplexing and beamforming, array-based systems have supported new verticals and applications since 3GPP Release 7 (Dreifuerst and Heath 2023b). Beam management is an important component of expanding MIMO to massive MIMO and extreme MIMO (X-MIMO) (Holma et al. 2021). Beam management employs codebook-based strategies for beam training with hybrid arrays (Heath et al. 2016; Hamid et al. 2023) and configurable feedback (Dreifuerst and Heath 2023b; Giordani et al. 2019) to support large-scale arrays. While beam management enables improved flexibility and higher dimensional MIMO arrays, the performance is dependent on the beam codebooks employed throughout the process. Yet optimizing codebooks for beam management is challenging because each codebook is dependent on the other system processes. As a result of this lack of optimization, massive MIMO, and even multi-user MIMO deployments are still limited in 5G (Dreifuerst and Heath 2024a).

Beam management, from a system level, is a series of operations enabling initial access, beam refinement, and feedback that ultimately lead to data transmission (Dreifuerst and Heath 2024a; Giordani et al. 2019). Beam tracking and other beam search reduction techniques can also be understood within the overarching beam management framework. The primary operations can be referenced by the codebooks and 3GPP signals transmitted during the respective stages. First, user equipment (UE) joins and synchronizes with the network using a pre-defined synchronization signal block (SSB) sequence that is beamformed using a beamforming codeword from a codebook stylized as the *SSB codebook* (Dreifuerst et al. 2022). Broadly speaking, this beamforming is intended to ensure all users can accurately synchronize, requires a small amount of time-frequency resources, and trades off beamforming gain for broad coverage. Second, beam refinement requires a large amount of resources to enable channel estimation on top of a fine-grained beam search using another beamforming codebook. Channel state information reference signals (CSI-RS) are transmitted using a subset of the *CSI-RS codebook* based on the user-reported SSB feedback. The beamformed channel represents a compressed channel and allows for reduced-rank feedback and increased coherence time (V. Va et al. 2017). Third, the *feedback codebook* is used to quantize the beamformed channel estimates to reduce the overhead of feedback signaling with large arrays. The feedback codebook is configurable and standardized in 5G, with the most relevant format for MU-MIMO (type-II) defined as an oversampled DFT codebook (Qin and Yin 2023). Using the combination of coarse and refined beam training as well as quantized feedback, a serving cell can design a set of precoders to serve the UEs for data transmission.

Although beam management has been standardized since 3GPP Release 15 (Dreifuerst and Heath 2023b), new research for optimizing beam management to improve the performance in 6G is underway (Lin 2023). At the forefront of this innovation are extensions and improvements for beam management including support for 128 ports (3GPP 2023b). Configuring, optimizing, and updating codebooks is challenging, though, due to codebook interrelationships between codebooks, user distributions, and realistic channel models. As a result, 3GPP has identified beam management and CSI as potential areas to integrate artificial intelligence/machine learning (AI/ML) (Lin 2023). These data-driven techniques present new possibilities for designing site-specific algorithms without explicitly modeling distributions within the data (Jiang et al. 2017; Zhang et al. 2019; Heng et al. 2021). Within the scope of AI/ML for beam management, we identify three directions of work relevant to our investigation: codebook design, beam training, and interference mitigation.

Prior work on codebook design: AI/ML is used in the design or implementation of beamforming codebooks (Shafin et al. 2020; Heng et al. 2022; Dreifuerst et al. 2022) or CSI feedback codebooks (Turan et al. 2024; Chen et al. 2023; Lee et al. 2023; Ma et al. 2023). Although beamforming and feedback codebooks are used for different roles, the techniques for applying AI/ML to these problems are very similar. For beamforming codebooks, hierarchical codebooks have been a primary focus and adopted for 5G since early proposals (Wang et al. 2009; Alkhateeb et al. 2014; Morozov et al. 2016). Prior work has considered supervised learning (Heng et al. 2022; Dreifuerst et al. 2022; Yang et al. 2023) and reinforcement learning (Shafin et al. 2020) styles of learning initial access codebooks. That work ((Heng et al. 2022; Dreifuerst et al. 2022; Yang et al. 2023; Shafin et al. 2020)) does not consider how the new beamforming codebooks impact beamformed channel feedback nor the achievable rate in MU-MIMO data transmission. Ignoring the system-level impact of codebook design is a reoccurring result in prior work on beam management where a subset of the process–either beam training or codebook design–is modified but overarching metrics such as throughput are not considered.

Codebook design using AI/ML has also been explored for feedback quantization and reconstruction (Wen et al. 2018; Chen et al. 2022; Lee et al. 2023; Turan et al. 2024; Chen et al. 2023; Xiao et al. 2023). Deep learning autoencoder models have been used to replace feedback codebooks (Wen et al. 2018; Chen et al. 2022) or as a denoising process to improve the received feedback (Lee et al. 2023; Turan et al. 2024; Chen et al. 2023; Xiao et al. 2023). Feedback autoencoders require the encoder and decoder of the model to be trained together and then shared throughout the network. While sharing a neural network may be reasonable for higher-level tasks, the size of the models and difficulty with standardizing these models make them ill-suited for practical deployments. Denoising algorithms that operate in a one-sided fashion (Lee et al. 2023; Turan et al. 2024; Chen et al. 2023) are advantageous as they can be employed at the BS and improve the uplink feedback, which typically operates at a low SNR due to uplink power constraints (Becirovic et al. 2022). The challenge with denoising algorithms is they tend to be site-specific, not generalize well, and are not applicable for dynamic beamformed channels as seen by hybrid arrays (Dreifuerst and Heath 2024a). Some work on generalization or meta-learning has been proposed (Xiao et al. 2023) though it still ignores the prior beam management steps that impact the effective channel and CSI distribution.

Prior work on beam training: AI/ML is also used in beam training, which is the process of determining the beam-codewords to use from the standard beamforming codebooks. Beam training has received significant attention toward the goal of maximizing the wireless link power with limited overhead (Yang et al. 2022; Sim et al. 2018; Wang et al. 2009; Choi et al. 2016; V. Va et al. 2018; Hussain and Michelusi 2022). Data-driven techniques have been used to integrate localization and sensing information (Choi et al. 2016), capture site-specific characteristics (Yang et al. 2023), or improve mobility robustness (V. Va et al. 2018; Rebato et al. 2018; Hussain and Michelusi 2022). Additional work on a generalized beam selection algorithm was proposed in (Vuckovic et al. 2023). These investigations are focused on mmWave bands where the environment is sparse. While sparse environments can challenge beam training, the sparsity also results in the beam search primarily focused on finding a rank-1 beamformer aligned with a single path. Lower bands such as sub-6GHz or upper-mid bands (Holma et al. 2021) exhibit richer scattering that makes the beam training process depend on coherent combining over multiple paths and limiting multi-user interference instead of simply finding a main beamforming direction. Instead, these works (Yang et al. 2022; Sim et al. 2018; Wang et al. 2009; Choi et al. 2016; V. Va et al. 2018; Hussain and Michelusi 2022) are primarily targeting analog beamforming and do not consider MIMO formats. This is a significant limitation in mobile broadband, which has supported SU-MIMO and digital systems since 3GPP release 7 (Dreifuerst and Heath 2023b).

Prior work on codebook-based interference mitigation: AI/ML interference-aware beam management is a combination of codebook design and beam training with additional consideration for how overlapping beams cause interference in adjacent sectors and cells. In one line of work, a conventional beam management interference cancellation algorithm was converted to a deep learning approximate to reduce complexity (Zhou et al. 2019). A direct AI/ML algorithm for interference-aware codebooks was also recently proposed (Zhang and Alkhateeb 2024). These algorithms ((Zhou et al. 2019; Zhang and Alkhateeb 2024)), though, are trained on data from mmWave bands and therefore benefit from the sparse channel nature while also only considering simplified SU-MISO systems. There is a lack of research investigating data-driven codebook algorithms that operate in rich scattering environments and in MU-MIMO settings.

Prior work on beamspace codebook design: Our prior work introduced the concepts of beamspace codebook learning and hierarchical codebook learning for X-MIMO arrays (Dreifuerst and Heath 2024a, 2023a). We proposed a learning algorithm that produced highly effective beam training codebooks but only provided a modest improvement in the overall spectral efficiency. In our present work, we now extend the end-to-end training to encapsulate the achievable spectral efficiency, thereby producing codebooks that support full rank MIMO connections, while also mitigating inter-cell interference. Contributions of this paper:

- We describe a complete beam management process for a multi-cell, hybrid, X-MIMO system. The beam process includes inter- and intra-cell interference during beam training and data transmission to characterize the impact of interference with dynamic codebooks. By capturing and modeling each stage from initial access through network spectral efficiency, we can obtain insights into the net effects of advanced codebook algorithms in realistic settings.

- We propose an end-to-end learning strategy and neural processing architecture "Network Beamspace Learning" (NBL) to design SSB and CSI-RS codebooks for multiple base stations serving overlapping regions. The strategy is based on representing beamformers in the beamspace and backpropagating gradients through the entire beam management model from initial access through achievable rate calculation. The end-to-end training strategy ensures that codebooks are learned to maximize the beamformed channel and mitigate interference that would reduce spectral efficiency. The NBL formulation only requires sharing the SSB feedback, yet significantly improves the overall system performance.

- We evaluate the capabilities of the beam management framework with X-MIMO using traditional and NBL codebooks across a range of possible deployments. In particular, we evaluate using different array geometries and codebook sizes in ray-traced environments. Ultimately, we find that MU-MIMO systems are susceptible to interference with extremely large arrays, but learned codebooks can mitigate the interference and achieve significantly higher cell spectral efficiency.

Notation: $\mathbf{A}$ is a matrix, $\mathbf{a}$ and $\{a[i]\}_{i=1}^{N}$ are column vectors and a, A denote scalars. $\mathbf{A}^T$, $\overline{\mathbf{A}}$, $\mathbf{A}^*$, and $\mathbf{A}^\dagger$ represent the transpose, conjugate, conjugate transpose, and psuedo-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 $\|\mathbf{a}\|_2 = \mathbf{a}^*\mathbf{a}$ for vectors and the Frobenius norm $\|\mathbf{A}\|_F = \sqrt{\mathrm{Tr}(\mathbf{A}\mathbf{A}^*)}$ for matrices. The operator $\mathbb{E}[\cdot]$ is used for the expectation of a

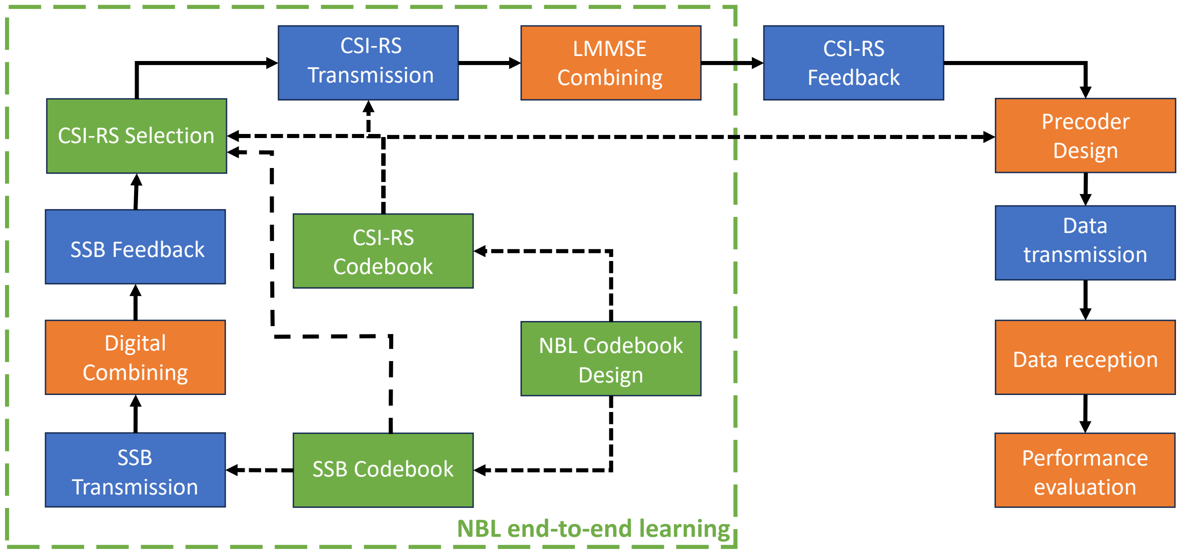

Figure 4.1: A block diagram of the system-level model and NBL codebook roles. The proposed solution modifies the green blocks, while the orange and blue blocks refer to assumed and standard operations respectively. Assumed operations are not standardized, but we make reasonable algorithmic decisions based on typical hardware, timing, and compute constraints. The end-to-end learning framework is shown by the green dashed box that encapsulates the codebook operations. The assumed operations outside of the end-to-end learning are not critical but allow for system-level evaluation.

random variable. Due to the notational complexity of MU-MIMO with OFDM, we will always use c and u to refer to a specific base station cell and UE, t as a specific time, k as a specific frequency resource, and $n_{\mathrm{t}}/n_{\mathrm{r}}$ to refer to a specific transmit or receive antenna. We use the superscript $(\cdot)^{\mathrm{RF}}$ to refer to the analog component of hybrid variables within the MIMO system.

The remainder of our paper is organized as follows. First, we introduce the multi-cell beam management model that includes the SSB and CSI-RS processes as depicted by the block diagram in Figure 4.1. We carefully define the transmission and reception of these signals as these expressions are directly integrated into the NBL backpropagation. Then, we describe a typical feedback and data transmission strategy that we use for evaluating the system-level performance of the proposed codebook algorithm in Sectoin 4.4. In Section 4.5 we describe NBL which includes beamspace preprocessing, a multi-stage neural network, and achievable-rate maximizing end-to-end learning. In the final sections, we present the simulation setup and a rigorous evaluation of the various codebooks, overall performance, and out-of-distribution (OOD) performance achieved by our proposed method.

4.3 Multi-cell beam management

In this section, we define the hierarchical beam training operation for a multi-cell scenario. Although beam training is standardized in 5G, we explicitly define each step of the beam-based operations because the multi-cell beam management process is integrated into the gradient updates of the learning algorithm in Section 4.5.

We consider a region that is served by $C = 3$ base stations (BS) or cells with overlapping coverage regions. Exterior cells are not simulated in this investigation, although extending the model and algorithm to include more BSs is straightforward with sufficient or dedicated compute nodes. Each BS is equipped with a planar array of size $N_{\mathrm{T}}^{\mathrm{RF}} = 2N_{\mathrm{X}}^{\mathrm{RF}} \times N_{\mathrm{Y}}^{\mathrm{RF}}$ with 2 horizontal elements of perpendicular polarization being colocated in the array. The antenna elements are connected in a fully connected hybrid format with b_{phase} bit phase shifters connecting to the $2N_{\mathrm{X}} \times N_{\mathrm{Y}}$ digital ports (Alouzi et al. 2023). We assume there are U UEs in the scene where U is a random variable that is not known to the network. Each UE is equipped with an N_{R} element fully digital uniform linear array enabling elevation beamforming. We assume a multi-cell MU-MIMO

OFDM channels between each antenna pair as

$$\mathbf{H} \in \mathbb{C}^{C \times U \times T \times K \times N_{\mathrm{R}} \times N_{\mathrm{T}}^{\mathrm{RF}}}. \tag{4.1}$$

The specific channel model is not critical as underlying distributions of the channel are learned from channel samples.

For a realized channel, we define a received signal model for a user u at time-frequency resource t, k with receiver combiner $\mathbf{W}_{u,t,k}^{*}$ and hybrid precoders $\mathbf{F}_{c,t,k}$, $\mathbf{F}_{c,t}^{\mathrm{RF}}$. We will assume timing and frequency synchronization have been performed using standardized reference signals. We also assume all users can be associated with a BS so that the U users can be partitioned $U = \sum_{c} U_c$ for ease of notation, although the user assignment may not be known prior to initial access or shared between cells. We define the desired signal, interference, and received signal as

$$\mathbf{D}_{u,t,k} = \mathbf{H}_{c,u,t,k}\mathbf{F}_{c,t}^{\mathrm{RF}}\mathbf{F}_{c,t,k}\mathbf{s}_{c,t,k} \tag{4.2}$$

$$\mathbf{I}_{c,u,t,k}^{\mathrm{intra}} = \sum_{c' \neq c}^{C} \mathbf{H}_{c',u,t,k}\mathbf{F}_{c',t}^{\mathrm{RF}}\mathbf{F}_{c',t,k}\mathbf{s}_{c',t,k} \tag{4.3}$$

$$\mathbf{y}_{u,t,k} = \mathbf{W}_{u,t,k}^{*}\left(\mathbf{D}_{u,t,k} + \mathbf{I}_{c,u,t,k}^{\mathrm{intra}} + \mathbf{N}_{u,t,k}\right). \tag{4.4}$$

The noise $\mathbf{N}_{u,t,k}$ is modeled as independent Gaussian random values and accounts for thermal noise and hardware noise figures. Note that $\mathbf{D}_{u,t,k}$ contains the desired signal for a given user and the inter-cell interference from streams of data for other users, while the intra-cell interference is contained in (4.3). The digital precoders $\mathbf{F}_{c,t,k} \in \mathbb{C}^{N_{\mathrm{T}} \times N_{\mathrm{R}}}$ are block diagonal where a subset of the digital array is used to serve and mitigate interference for each user $u \in U_c$ (Dreifuerst and Heath 2024a). The analog precoders $\mathbf{F}_{c,t}^{\mathrm{RF}} \in \mathbb{C}^{N_{\mathrm{T}}^{\mathrm{RF}} \times N_{\mathrm{T}}}$ are a matrix of frequency-flat beamformers used to direct the data streams.

The task of determining analog and digital precorders that maximize the network-wide spectral efficiency (SE) is challenging in the face of limited CSI and coordination. Beam management is one strategy that involves multiple stages of beam training and feedback to obtain CSI and configure the precoders. In particular, the beam training stages are used to obtain coarse and refined beamformers that make up the analog beamforming, while user feedback is used to configure the digital precoding (Giordani et al. 2019; Dreifuerst and Heath 2024a). The focus of our work here is on optimizing network performance by dynamically generating new beamforming codebooks.

For the duration of our paper, we assume a working knowledge of the beam management stages and will only briefly highlight the key equations. Readers are encouraged to review (Dreifuerst and Heath 2023b,a; Giordani et al. 2019; Dreifuerst and Heath 2024a) for an in-depth review of these processes in 5G. Given the focus on beam management for 6G (Lin 2023), we will assume the same three phases (SSB, CSI-RS, and feedback) as in 5G. We also assume UEs operate similarly to sub-6GHz 5G implementations, although such assumptions are just a simplification to focus on the BS operation. Much of the beam training and hybrid operations are not specified in the standards, so we will assume a formulation similar to classical hybrid implementations found in literature i.e. (Heath et al. 2016; Singh and Joshi 2021; Wu et al. 2018) and following the unified hybrid beam management system model from our prior work (Dreifuerst and Heath 2024a). One important difference is the multi-cell formulation, where we will assume the *nearby cells operate over the same bands and perform beam training on the same time-frequency resources.* This assumption is a worst-case scenario as nearby base stations could employ different bands for SSB and CSI-RS to mitigate multi-cell interference since the inception of beam management in 3GPP Release 15.

4.3.1 Synchronization signal block transmission

Initial access begins by transmitting a small number of broadcast beams with synchronization and basic network configuration information (Giordani et al. 2019; Dreifuerst and Heath 2023b). The process of transmitting a burst of SSBs is repeated periodically and allows new users to join the network. Historically, SSB codebooks were simple codebooks designed to form wide beams to cover large areas with limited beamforming gain (Raghavan et al. 2016). Directional beamforming, however, is not well-suited for multipath environments as experienced in sub-6GHz and frequency range (FR) 1 bands (Dreifuerst and Heath 2023a). With the upcoming increased support for X-MIMO arrays expected in 6G, new codebooks that operate in rich scattering environments and across array sizes are critical.

The SSB signals are transmitted on frequency resources K^{SSB} and specified timeslots T_i^{SSB} for each SSB i up to the total SSB size $L_{\max}$. The received SSB signal before combining can be simplified from (4.4) to a single-stream representation of the demodulation reference signals $s_{c,t,k}^{\mathrm{DMRS}} \in \mathbb{C}$ as

$$\mathbf{y}_{c,u,t,k}^{\mathrm{SSB}_i} = \frac{1}{\sqrt{KN_{\mathrm{T}}}}\mathbf{H}_{c,u,t,k}\mathbf{f}_{c,i}s_{c,t,k}^{\mathrm{DMRS}} \tag{4.5}$$

$$+\sum_{c'\neq c}^{C}\mathbf{H}_{c',u,t,k}\mathbf{f}_{c',i}s_{c',t,k}^{\mathrm{DMRS}} + \mathbf{n}_{u,t,k} \tag{4.6}$$

where $\mathbf{f}_{c,i} \in \mathbf{B}_c^{\mathrm{SSB}}$ is the analog beamformer controlled from a single digital port. Effectively, a digital beamformer could also be used to compensate for low-resolution phase shifters in the hybrid array, although the effects are almost imperceptible during rank-1 beamforming (Alouzi et al. 2023; Dreifuerst and Heath 2024a). We will assume the UE performs digital combining for each cell-specific reference signal to maximize the SSB

RSRP as

$$\mathrm{RSRP}_{c,i,u} = \sum_{k \in K^{\mathrm{SSB}}} \sum_{t \in T_i^{\mathrm{SSB}}} \left\| \mathbf{y}_{c,u,t,k}^{\mathrm{SSB}_i} \right\|^2. \tag{4.7}$$

In general, UEs could be configured for beam training or other combining strategies, but our focus is on the base station operations so we assume digital combining for simplicity.

After SSB reception, the UEs provide a short feedback packet that includes a selected base station, SSB index, and corresponding RSRP obtained during the selected SSB. In total, the aggregate feedback in the network is represented as

$$\mathbf{p} = \left\{ \max_i \mathrm{RSRP}_{c,i,u} \right\}_{u=1}^{U} \tag{4.8}$$

$$\mathbf{b}, \mathbf{m} = \left\{ \arg\max_{c,i} \mathrm{RSRP}_{c,i,u} \right\}_{u=1}^{U}. \tag{4.9}$$

The vectors $\mathbf{m}$ and $\mathbf{b}$ correspond to the best cell-beam index (4.9) and $\mathbf{p}$ is the corresponding RSRP (4.8) for all U users. Note that $\mathbf{p}$ would be quantized in realistic systems, although the quantization resolution is $\leq$ 1dB so we neglect the RSRP quantization for this investigation. We will formulate the problem around a weakly centralized controller that obtains the slowly-updated feedback from (4.8)-(4.9). In general, however, each cell may only have access to the information of users for its cell, especially for the CSI-RS step following SSB feedback. The cell-specific feedback is denoted

$$\mathbf{p}^{(c)} = \left\{ \max_i \mathrm{RSRP}_{c,i,u} \right\}_{u \in U_c} \tag{4.10}$$

$$\mathbf{m}^{(c)}, \mathbf{b}^{(c)} = \left\{ \arg\max_{c,i} \mathrm{RSRP}_{c,i,u} \right\}_{u \in U_c}. \tag{4.11}$$

After the SSB feedback, each BS can transmit CSI-RS to obtain more accurate beam alignment and channel estimates in the form of PMI.

4.3.2 CSI-RS transmission

Refined beam training is performed by transmitting CSI-RS using a hierarchical codebook strategy following SSB feedback. CSI-RS occupy significantly larger resources and enable pilot-based channel estimation. Typical CSI-RS deployments begin with a large CSI-RS codebook ($\gg L_{\max}$) and then perform a subset selection based on the SSB feedback (Dreifuerst and Heath 2023a). We define $\mathbf{B}^{\text{sub}}$ as the active CSI-RS codebook subset and employ a selection algorithm similar to (Dreifuerst and Heath 2024a). The important difference is instead of searching over each vector, the search is performed over precoder matrices, and the most correlated vectors from each precoder are compared for subset selection. The benefit to precoder-based selection is that the goal is no longer to achieve a set of rank-1 beamformers and then spatially multiplex with them independently; instead, the codebook contains sets of precoders that are likely to support multi-stream communication. In particular, we modify (Dreifuerst and Heath 2024a, equation (8)) to instead be

$$\mathbf{C} = \max_{s} \left\langle \mathbf{B}_{c,i}^{\text{SSB}}, \mathbf{B}_{c,j,s}^{\text{CSI-RS}} \right\rangle \quad \forall i, j. \tag{4.12}$$

From there, each step is essentially identical with the selected codewords now corresponding to matrices instead of grouping the beamforming vectors in the following step. The CSI-RS subset selection process is repeated for each cell independently using the cell-specific feedback $\mathbf{p}^{(c)}, \mathbf{m}^{(c)}, \mathbf{b}^{(c)}$. The result is a subset of N_{CSI} CSI-RS beams for each base station that is proportionally selected to have at least one beamformer that is highly correlated with the active SSB beams reported.

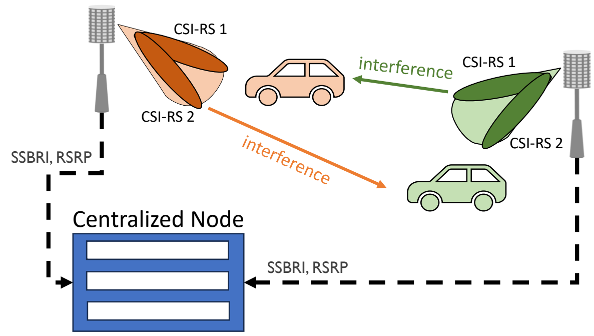

Figure 4.2: Users are allocated to the cells during SSB transmission. During CSI-RS transmission, only interference during a UE's selected CSI-RS causes interference. Therefore, codebooks should be designed so that the first green CSI-RS does not interfere strongly with the orange UE and vice versa.

The cells begin the CSI-RS transmission process by transmitting training pilot symbols $\mathbf{s}^{\mathrm{tr}}_{c,t,k} \in \mathbb{C}^{B_{\mathrm{g}}\times 1}$ with the active subset of beamformers $\mathbf{B}^{\mathrm{sub}}_{c}$. The received CSI-RS signal for each UE before combining is then

$$\mathbf{y}^{\text{CSI-RS}_i}_{u,t,k} = \mathbf{H}_{\mathbf{b}_u,u,t,k}\mathbf{B}^{\mathrm{sub}}_{\mathbf{b}_u,i}\mathbf{s}^{\mathrm{tr}}_{\mathbf{b}_u,t,k} \tag{4.13}$$

$$+ \sum_{b'\neq \mathbf{b}_u} \mathbf{H}_{b',u,t,k}\mathbf{B}^{\mathrm{sub}}_{b',i}\mathbf{s}^{\mathrm{tr}}_{b',t,k} + \mathbf{n}_{u,t,k}. \tag{4.14}$$

Importantly, the users are now associated with a given base station $\mathbf{b}_u$ determined during the SSB state. Therefore, all other cells create interference as a result of their CSI-RS transmissions in the same time-frequency blocks as shown in Figure 4.2. The UE, however, can utilize its antenna array to combine the intended signal and mitigate interference, for example, with a linear minimum mean squared error (LMMSE) receive combiner

(Heath Jr. and Lozano 2018, Section 9.4) resulting in

$$\mathbf{R}_{u,t,k} = \left(\sum_{c}^{C} \mathbf{H}_{c,u,t,k} \mathbf{B}_{c,i}^{\mathrm{sub}} (\mathbf{H}_{c,u,t,k} \mathbf{B}_{c,i}^{\mathrm{sub}})^{*} \right) \tag{4.15}$$

$$\mathbf{V}_{i,u,t,k} \triangleq \mathbf{H}_{b_u,u,t,k} \mathbf{B}_{b_u,i}^{\mathrm{sub}} \tag{4.16}$$

$$\mathbf{W}_{u,t,k}^{\text{CSI-RS}_i} = \mathbf{R}_{u,t,k}^{-1} \mathbf{V}_{i,u,t,k} \tag{4.17}$$

$$s_{u,t,k,n_{\mathrm{r}}}^{\text{CSI-RS}_i} = \frac{(\mathbf{V}_{i,u,t,k})_{n_{\mathrm{r}}}^{*} \mathbf{R}_{u,t,k}^{-1} (\mathbf{V}_{i,u,t,k})_{n_{\mathrm{r}}}}{1 - (\mathbf{V}_{i,u,t,k})_{n_{\mathrm{r}}}^{*} \mathbf{R}_{u,t,k}^{-1} (\mathbf{V}_{i,u,t,k})_{n_{\mathrm{r}}}}. \tag{4.18}$$

Equation (4.18) calculates the signal-to-interference-noise ratio (SINR) based on receiver combiner $\mathbf{W}_{u,t,k}^{\text{CSI-RS}_i}$ in (4.17). The SINR is a configurable reporting metric in 5G and is measured using zero-power and non-zero-power CSI-RS (Dreifuerst and Heath 2023b).

Finally, the UE selects the CSI-RS resource that would maximize the rate based on a configurable metric estimated from the pilot symbols. The two standardized metrics are SINR or RSRP, although SINR is preferred for interference management. Mathematically the selection process is based on the average over the time-frequency resource $(T^{\text{CSI-RS}_i}, K^{\text{CSI-RS}_i})$

$$\mathrm{SE}_{u}^{\text{CSI-RS}_i} = \sum_{n_{\mathrm{r}}=1}^{N_{\mathrm{R}}} \log_2 \left(1 + \frac{1}{T^{\text{CSI-RS}_i} K^{\text{CSI-RS}_i}} \sum_{t,k} s_{u,t,k,n_{\mathrm{r}}}^{\text{CSI-RS}_i} \right) \tag{4.19}$$

$$\widehat{\mathbf{i}}_{c,u} = \left\{ \arg\max_{i} \mathrm{SE}_{u}^{\text{CSI-RS}_i} \right\}_{u \in U^c}. \tag{4.20}$$

The analog beam selections, $\widehat{\mathbf{i}}_{c,u}$ correspond to the chosen analog precoders that will be used by the base station during hybrid data transmission because all of the UE feedback is based on the beamformed channels, $\mathbf{H}_{b_u,u,t,k} \mathbf{B}_{b_u,\widehat{i}_{b_u,u}}^{\mathrm{sub}}$.

Compared to the SSB model (4.6), the CSI-RS received model is multi-layer due to the need for sounding each of the B_{g} ports for each user.

4.4 System model

In this section, we define a typical implementation of feedback and data transmission using the CSI available through beam management. While the previous section is integrated within the NBL learning process, this section can be understood as the "system model" surrounding our algorithm that allows for evaluating the proposed algorithm within a realistic system. The system model can be thought of as modular within wireless systems, where the algorithms for feedback or data transmission can be replaced without impacting the proposed solution.

4.4.1 Feedback

The feedback digitally encodes and reports the channel state information obtained during CSI-RS. In this section, we motivate the focus on SINR in the previous section as it relates to feedback, define relevant variables for future steps, and discuss assumptions made for the feedback. For additional discussion on feedback, we refer readers to (Dreifuerst and Heath 2023a; Qin and Yin 2023).

The SINR (4.18) impacts the channel estimation efficacy and is useful for determining the modulation and coding scheme (MCS) for data transmission. The beamformed channel estimate is defined for a channel estimation algorithm $f(\cdot)$ as

$$\widehat{\mathbf{H}}_{u,t,k} \triangleq f(\mathbf{y}_u^{\text{CSI-RS}_i}, \mathbf{s}_{c,t,k}^{\text{tr}}) \tag{4.21}$$

For ease of notation, we define $\widehat{\mathbf{H}}_{u,t,k}$ as the beamformed channel estimate–rather than the physical channel estimate–with the assigned base station corresponding to $\mathbf{H}_{b_u,u,t,k}\mathbf{F}_{b_u,i_{b_u,u}}$. There is no standardized channel estimation algorithm. Given the limited CSI-RS response timing and power budget of mobile UEs, an LMMSE method is assumed herein. The resulting channel estimate at the UE is frequency selective over a set of frequency subbands $b \in \{1, ...\mathrm{S_B}\}$ and carried out in a similar manner to (Dreifuerst and Heath 2023a). For ease of notation, we focus on a single subband although employ multiple subbands in the simulation results in Section 4.7. In the final step, the channel estimates are quantized according to a feedback codebook and transmitted to the BS (Dreifuerst and Heath 2023b; Qin and Yin 2023).

The signal-to-interference-noise ratio (SINR) presented in (4.18) provides a valuable secondary metric for the beam management efficacy as it ultimately determines the channel estimation performance as a result of beam alignment with the desired serving cell and interference from nearby cells. In Section 4.5 we describe how we use the SINR and achievable spectral efficiency (4.19) resulting from the entire beam management process in the end-to-end learning framework.

4.4.2 Data transmission

Following the first three stages of beam management (SSB, CSI-RS, and feedback), the network can configure and transmit data using the combined analog precoders $\mathbf{F}_{b_u,\hat{i}_{b_u,u}}$ and corresponding feedback $\widehat{\mathbf{H}}_u$. In particular, we formulate the precoder design algorithm in the same way as traditional hybrid formulations to make use of well-known techniques (Heath Jr. and Lozano 2018; Heath et al. 2016; Singh and Joshi 2021). Conceptually, the analog beams chosen from the CSI-RS codebook provide beamforming gain, while a regularized zero-forcing precoder (RZF) is used to configure the digital precoding to mitigate intracell interference. Furthermore, each cell does not have feedback for users that are associated with other cells, so each cell operates independently of the other cells at this stage. Therefore, the following steps directly follow the single-cell RZF formulation from (Dreifuerst and Heath 2024a) with the key equations included here for reference.

RZF precoding is based on maximizing the signal-to-leakage noise ratio, thereby allowing for some inter-stream interference to improve low-SNR performance (Heath Jr. and Lozano 2018, section 9.9). Without loss of generality, we assume the users are ordered so that the first U_c^{a} users are actively selected for joint transmission according to some scheduling algorithm. Then the analog precoding is defined as

$$\mathbf{F}_c^{\mathrm{RF}} = [\mathbf{F}_{c,\hat{i}_0}, \mathbf{F}_{c,\hat{i}_1}, ..., \mathbf{F}_{c,\hat{i}_{U_c^{\mathrm{a}}}}]. \tag{4.22}$$

The digital precoders are determined for each user and associated in a block diagonal form based on the rank of the CSI-RS B_{g}

$$\mathbf{F}_{c,u,t,k} = \frac{(\sum_{i=0}^{U_c^{\mathrm{a}}-1} \widehat{\mathbf{H}}_{i,t,k}^* \widehat{\mathbf{H}}_{i,t,k} + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^2])^{-1}\widehat{\mathbf{H}}_{u,t,k}^*}{\left\|(\sum_{i=0}^{U_c^{\mathrm{a}}-1} \widehat{\mathbf{H}}_{i,t,k}^* \widehat{\mathbf{H}}_{i,t,k} + UN_{\mathrm{T}}\mathbb{E}[\mathbf{N}_{u,t,k}^2])^{-1}\widehat{\mathbf{H}}_{u,t,k}^*\right\|} \tag{4.23}$$

$$\mathbf{F}_{c,t,k} = \begin{bmatrix} \mathbf{F}_{c,0,t,k} & 0 & ... & 0 \\ 0 & \mathbf{F}_{c,1,t,k} & 0 & ... \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & ... & \mathbf{F}_{c,U_c^{\mathrm{a}},t,k} \end{bmatrix}. \tag{4.24}$$

With the precoding defined, we can finalize the steps necessary to reach the network metric: effective sum spectral efficiency (ESSE). We assume the same LMMSE receive combining strategy by the UEs as defined in (4.17), with the slight modification for the data precoders instead of the CSI-RS precoders. Then the per-user SINR is direct copy of (4.18) with the replaced variables. Finally, the network-wide sum spectral efficiency after removing the time-frequency resources used for beam management $(T_{\mathrm{BM}}, K_{\mathrm{BM}})$ is

$$\mathrm{ESSE} = \sum_{u=1}^{U} \sum_{\substack{t \notin T_{\mathrm{BM}} \\ k \notin K_{\mathrm{BM}}}} \sum_{r=1}^{R} \log_2(1 + s_{u,t,k,r}). \tag{4.25}$$

Note that the beam management process up through data transmission is repeated regularly throughout the wireless network. During each period, the codebooks can be retained and new codebooks can be generated to best serve the specific site characteristics including user dynamics and environment characteristics. The periodic nature of beam management is the basis of our proposed algorithm, NBL, which employs deep learning to generate codebooks between beam management periods based on the limited feedback and prior codebook information. By employing deep learning to generate codebooks, as opposed to channel estimation or precoder design, the system can circumvent many of the timing and compute constraints of deep learning when it comes to wireless communications. In addition, by explicitly modeling the beam management steps, a neural network can be trained with gradients calculated through the entire process in an end-to-end style (Felix et al. 2018).

4.5 Network beamspace learning

In this paper, we propose a gradient-complete multi-cell MU-MIMO codebook optimization strategy. We previously found that neural codebook design based on beamspace representations is a powerful tool for codebook design (Dreifuerst and Heath 2023a). Training the architecture to maximize beam training power was also effective (Dreifuerst and Heath 2024a), but ultimately focused on an intermediate goal–beam training power–rather than the network goal of spectral efficiency maximization. Conceptually, in this

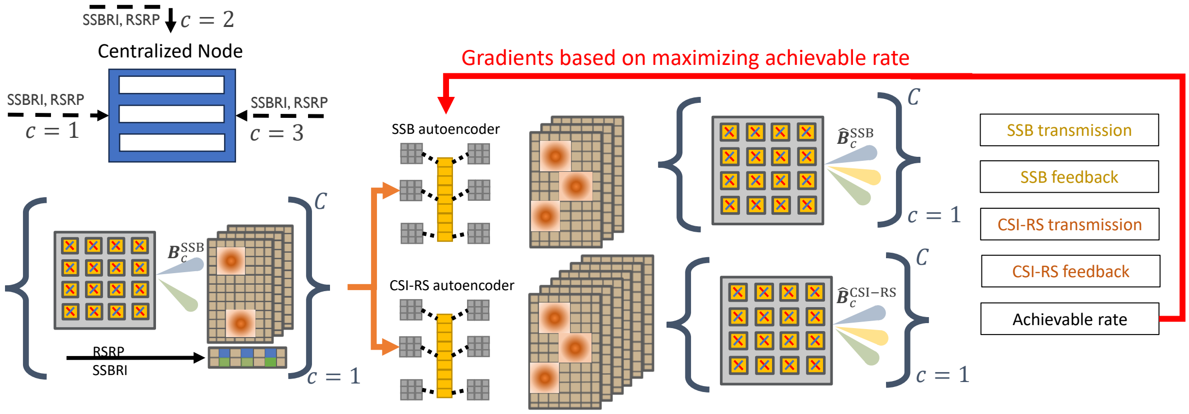

Figure 4.3: A visual representation of the NBL inference and training within the network. SSB feedback is shared between cells to jointly generate SSB and CSI-RS codebooks. Gradients are backpropagated through the entire beam management system to maximize the achievable rate during CSI-RS.

paper, we reframe the learning objective to maximize the sum spectral efficiency in the presence of multi-user and multi-cell interference. In this section, we define the end-to-end neural architecture as shown in Figure 4.3, beginning with the beamspace preprocessing and including up through codebook generation. We define three important aspects of the neural architecture design process that fundamentally aid the NBL algorithm.

Design aspect 1: *Encapsulating the beam management framework within the end-to-end learning produces NBL codebooks that maximize system-level performance.* Learning SE maximizing codebooks requires more than simply training the neural network to predict a specific set of beamformers in a typical supervised fashion. Instead, we capture the beam management process as described in Section 4.3 in the backpropagation step. Then the neural network is trained to maximize the achievable rate during CSI-RS based on (4.19). Ultimately, because of the reliance on the CSI-RS codebook for analog beamforming in the data transmission, the codebooks used during beam training should support high data rates rather than simply rank-1 transmissions.

The first step in the neural processing pipeline is converting the previous SSB Codebook into a beamspace representation (Dreifuerst and Heath 2023a, 2024a). The beamspace representation provides a consistent learning paradigm where the SSB feedback corresponds to an "image" of the beamforming magnitude over a range of azimuth and elevation directions. In addition to setting up the learning paradigm, the beamspace representation also allows the neural architecture to generalize to any antenna size or number of users by separating those dimensions from the inputs and outputs of the neural network. To be precise, neural networks are often limited by requiring consistent dimensions for inputs and outputs but we circumvent this through the beamspace representation which allows for arbitrarily sizing the input "image" to sampling dimensions $N_{\mathrm{X,O}} \times N_{\mathrm{Y,O}}$ in the azimuth and elevation directions for each beamformer. The only requirement is that $N_{\mathrm{X,O}} \geq N_{\mathrm{X}}$ and $N_{\mathrm{Y,O}} \geq N_{\mathrm{Y}}$ to ensure the beamspace is not aliased. The beamspace transformation can be performed at each cell and shared with local cells or a small centralized node during each SSB period ($\approx$ 20ms) to form a *weakly-centralized* setting (see Figure 4.3). The centralization or SSB sharing is motivated by the long SSB periods as well as only sharing the SSB feedback between cells, rather than large packets of information.

We define the beamspace representation for each cell c as $\mathbf{O}_c^{\mathrm{BSC}} \in \mathbb{C}^{(N_{\mathrm{X,O}}+1)\times(N_{\mathrm{Y,O}}+1)\times(L_{\max})}$ formulated in the same way as (Dreifuerst and Heath 2024a). First, transformation matrices $\mathbf{U}_{N_1,N_2}$ are calculated from the array responses $\mathbf{a}_N(\theta)$ (Heath et al. 2016) for antenna dimension N_1 to beamspace sampling dimension N_2

$$\boldsymbol{\theta}_{N_2} = \frac{1}{\pi}[0, 1, ..., N_2 - 1]^T \tag{4.26}$$

$$\mathbf{U}_{N_1,N_2} \triangleq [\mathbf{a}_{N_1}(\boldsymbol{\theta}_0), ...\mathbf{a}_{N_1}(\boldsymbol{\theta}_{N_2-1})]. \tag{4.27}$$

Then the beamspace representation is

$$\widehat{\mathbf{O}^{\mathrm{BSC}}}_{c,i} = \mathbf{U}^{*}_{N_{\mathrm{X}}^{\mathrm{RF}},N_{\mathrm{X,O}}}\mathbf{F}^{\mathrm{SSB}}_{c,i}\mathbf{U}_{N_{\mathrm{Y}}^{\mathrm{RF}},N_{\mathrm{Y,O}}} \quad \forall c,i. \tag{4.28}$$

$$\mathbf{O}_{c}^{\mathrm{BSC}} = \left\{\left[\widehat{\mathbf{O}^{\mathrm{BSC}}}_{c,i}, \sum \mathbf{1}_{\mathbf{m}=i}, \mathbf{1}^{T}_{\mathbf{m}=i}\mathbf{p}\right]\right\}_{i=0}^{L_{\max}-1}. \tag{4.29}$$

Note that in a centralized scenario, the codebook generation is carried out at a shared compute node and shared with each cell which in turn share the SSB feedback as shown in Figure 4.3. In the event of a fully disaggregated scenario, the interfering cells cannot be jointly optimized, the nearby cell's beamspace representations are not shared, and interference instead is treated as noise during the gradient updates.

Design aspect 2: *Beamspace representations enable generalization in dynamic codebook learning algorirthms.* The beamspace preprocessing forms a consistent basis that abstracts the carrier frequency and antenna geometry from the neural network. The abstraction is performed by the beamspace conversion that incorporates the array response, which is a function of the operating frequency and geometry. The beamspace representation allows for the same neural network to be used in a range of deployment strategies for network operators (Dreifuerst and Heath 2024a). An evaluation of the generalization capabilities is presented in Section 4.7.

After computing the beamspaces for each cell, the concatenated $\mathbf{O}^{\mathrm{BSC}}$ is fed into a $3+3$ layer convolutional autoencoder architecture with $(128, 128, 256)$ filters and 2×2 max pooling between layers. In contrast to our prior work (Dreifuerst and Heath 2024a), the final 3 layers perform an inverse convolution (typically termed convolution-transpose in deep learning frameworks) instead of using fully connected layers due to the large dimensional outputs which would require significant hardware resources (more than a single GPU) to parameterize in a fully connected architecture. By employing parameter sharing during the 3 encoding and 3 decoding layers, the network only requires ≈ 2 million parameters for the joint SSB and CSI-RS prediction for all cells.

The output of the neural network is a predicted beamspace which is then converted back to the complex codebook entries for the next beam training period. The deconversion process is performed in the same way as the conversion process using the matrix inverses of (4.28). With the new codebooks $\widehat{\mathbf{F}}^{\text{SSB}}, \widehat{\mathbf{F}}^{\text{CSI-RS}}$ we can introduce the end-to-end learning strategy for rate-maximized codebook learning. Ultimately, the end-to-end strategy allows for learning complex and arbitrary codebooks that minimize the loss in achievable rate from a perfect CSI, SU-MIMO, interference-free setting.

First, the target "labels", which are the maximum achievable spectral efficiencies, are computed from the channel assuming perfect CSI and SU-MIMO data transmission. The maximum achievable SE is calculated from the singular value decomposition (SVD) for a user u with connected BS c as

$$\mathbf{U}_u, \mathbf{S}_u, \mathbf{V}_u^* = \mathbf{H}_{c,u,t,k} \tag{4.30}$$

$$r_u = \sum_{n_\text{r}=1}^{N_\text{R}} \log_2 \left(1 + \frac{S_{u,n_\text{r}}^2}{\sigma^2} \right). \tag{4.31}$$

For ease of notation, we focus on a single time-frequency resource t, k, although the actual implementation is performed over an applicable set of resources. In a single-cell scenario, the system could then calculate the achievable spectral efficiency resulting from serving all active users with different analog precoder selections. When extending this to the multi-cell setting, however, it is pessimistic to assume that every analog precoder $\mathbf{B}^{\text{sub}}$ from every BS is active and causing interference. Furthermore, directly assuming all precoders cause interference ignores the timing aspects of beam management, where interference is acceptable during beams that a user does not end up selecting. In essence, the beams from nearby cells should cover non-overlapping regions at a given time. We capture this property by assuming the interference only occurs from the beams used during the CSI-RS of the selected CSI-RS $\widehat{\mathbf{i}}_{c,u}$ as in (4.18). Mathematically the achievable spectral efficiency can therefore be calculated as (4.19).

Design aspect 3: *Incorporating realistic time-frequency usage within the training allows the algorithm to coordinate beamforming between interfering cells.* Interference during CSI-RS causes channel estimation errors and CSI-RS beams may be allocated by users for the nearby cells. Assuming interference only during the user-selected CSI-RS resource frees the algorithm to learn analog precoders that do not interfere strongly during the same CSI-RS to improve channel estimation and subsequent data transmission. In a non-centralized format, the neural network would be limited to simply maximizing the signal strength, but this does not necessarily improve the network ESSE (see Figure 4.13 of Section 4.7). Then the overall loss function is the mean-squared-error (MSE) of the target and predicted achievable rates

$$\mathcal{L}(\mathbf{r}, \widehat{\mathbf{r}}) = \frac{1}{U} \sum_{u} (r_u - \widehat{r}_u)^2. \tag{4.32}$$

In an ideal scenario, the loss would approach 0 with perfect CSI and no noise. Of course, there is always non-zero noise and interference thereby causing the loss function to be strictly positive. Gradient calculation is backpropagated throughout the entire system model for the selected SSB and CSI-RS codewords.

4.6 Raytraced beam management dataset

System-level simulations that capture the entire beam management cycle are critical for evaluating the ultimate network effects of the proposed NBL codebooks. Simulation can be broken down into two steps: spatially consistent channel generation and system modeling. We employ Sionna (Hoydis et al. 2022), a raytracing channel simulator, that generates an offline dataset of user channels to sample from. In our results, we define a default scenario of a raytraced environment around the Frauenkirche in Munich, Germany operating at 10 GHz with three base stations 40m tall each equipped with a planar array $N_{\mathrm{X}} = 16$, $N_{\mathrm{Y}} = 16$, and 32 digital ports facing toward the serving region. We also consider a scenario with a smaller antenna array, $N_{\mathrm{X}} = 8$, $N_{\mathrm{Y}} = 8$, and another possible deployment scenario with the system operating in Paris, France at 20 GHz from a height of 27m. These scenarios, shown in Figures 4.4a-4.4b, allow us to consider urban environments in the low- and high-bands of the proposed upper-mid band (Holma et al. 2021) as well as different array geometries and performance in the face of OOD data. Users are assumed to have $N_{\mathrm{R}} = 4$ antennas and scattered outdoors throughout the scene at heights of 0.5-1.5m above ground. The channels are sampled every millisecond over a 100MHz bandwidth comprised of 270 resource blocks with 30kHz subcarrier spacing.

After generating channels for 1000 mobile UEs, a training dataset of beamspaces, channels, and target spectral efficiencies can be built up. For each dataset sample, a random number of active users is selected $U \sim \mathcal{U}[8, 20]$ and drawn from the channel dataset. Of those users, each is assigned a 20% probability of being considered a "new" user, and feedback of these users is not included in the subsequent beamspace calculation. Non-new users are associated with base stations based on the SSB feedback $\mathbf{b}$ determined from the prior SSB codebook, which is initialized with a DFT codebook. The beamspace

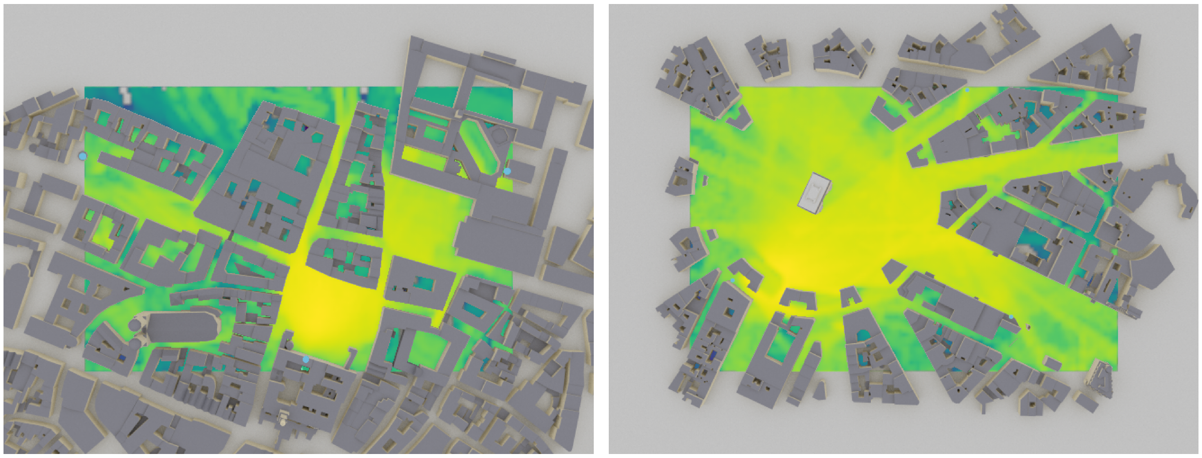

(a) Scene 1 (b) Scene 2

Figure 4.4: A comparison of the powermaps in two environments with different base station heights and carrier frequencies. Scene 1 corresponds to Munich, Germany operating at 10GHz and base stations mounted at 40m tall and 450m intersite distance. The OOD environment in scene 2 is Paris, France operating at 20 GHz, with 27m high serving cells and 370m intersite distance. A low resolution power map is rendered on each scene in the overlapping region from the perspective of base station 1 in each of the scenes.

is also generated from the prior SSB codebook as well as the per-cell SSB feedback. The SEs are calculated according to (4.31) and the triplet (beamspace, channels, and SE) are stored for the offline training. During training, the new SSB and CSI-RS codebooks are generated from the beamspace, SSB and CSI-RS feedback is simulated for the next beam training period, and the models are updated based on gradients calculated from the end-to-end loss (4.32). Gradient updates are performed with an Adam optimization and learning rate of 10^{-4} from a training dataset with a validation dataset used for determining and comparing model performance. All simulation results are calculated from a new, unseen test dataset. Overall, the network is trained with $20,000$ samples in the training datasets corresponding to different array sizes, environments, and carrier frequencies.

During model evaluation new channels are generated, the AI model is used to predict the codebooks, and the data transmission performance is calculated according to a greedy scheduler with RZF precoding (4.22)-(4.25). In particular, the base stations estimate the achievable sum SE for every combination of users from the coarse CSI reported and select the highest combination, i.e. there is no cooperation outside of the codebook generation. The ESSE (4.25) is then calculated from the true CSI for evaluating the overall performance of the proposed system. The resulting RSRP, SINR, and ESSE metrics are used for comparison in Section 4.7. DFT codebooks are used for comparison where the SSB codebook and CSI-RS codebook are wide and narrow DFT codebooks. In the DFT codebooks, the least commonly used beams are not considered. For example, beams angled upward are never used as these would never be selected by users scattered at ground level. The DFT selection provides us with a reasonable comparison with traditional codebooks while still incorporating a small degree of site-specific codebooks.

4.7 Simulation results

In this section, we present a careful evaluation of the proposed NBL algorithm. Neural networks are capable of function matching in a wide range of tasks, but also susceptible to overfitting (Vuckovic et al. 2023; Zhang et al. 2019). We consider scenarios where the training data and testing data come from different distributions as a result of different physical array geometries, deployment locations, and operating frequencies. The first set of results focuses on a site-specific scenario to evaluate the best-case performance; out-of-distribution settings are tested in the later subsections. Unless specified, we assume $L_{\max} = 16$, $N_{\mathrm{CSI}} = 16$, type-II PMI codebooks with $L_{\mathrm{CSI}} = 4$ (Dreifuerst and Heath 2023b; Qin and Yin 2023) and all results refer to Scene 1 from Figure 4.4.

4.7.1 SSB RSRP

The first evaluation focuses on the RSRP performance of the proposed NBL algorithm. The RSRP is evaluated using the SSB codebooks. There are two goals: 1) To achieve a sufficiently strong connection so that UEs are capable of accurate synchronization and 2) That the subsequent CSI-RS subset $\mathbf{B}^{\text{sub}}$ achieves a beamformed channel with high SINR for channel estimation and MU-MIMO SE. The first goal is directly quantified by the reported RSRP, with the goal typically being that all users are above the sensitivity, which is conservatively ≈ -120dBm. We show the empirical CDF of the reported RSRP in Figure 4.5 with different numbers of SSB beams $L_{\max}$. It can be seen that the DFT codebook produces twice as many users below the sensitivity and more than 5dB separates the bottom 20% of users. We can also see that NBL codebooks achieve better performance with the smallest codebooks $L_{\max} = 8$ than DFT codebooks ever reach as a result of the rich scattering environment. Current 5G implementations are limited to $L_{\max} = 8$ for sub-6GHz bands and $L_{\max} = 64$ in mmWave bands, so ML codebooks can help alleviate beam training overhead with X-MIMO arrays in 6G.

We also show a comparison of the per-user difference in RSRP between the two codebook methods in Figure 4.6. While the ECDF shows large-scale differences, it is also important to characterize if there are negative effects of the NBL codebooks. There are significant improvements as a result of the NBL codebooks with 93% of users increasing by more than 3dB. Yet, a small amount ($\leq 2\%$) of users experience a worse RSRP compared to DFT codebooks. Given the slight deficit among a small percentage of the population, this further validates the improvements in the SSB codebook performance as a result of learned codebooks. We also compare the load balancing of codebooks in Table 4.1. Overall there is not a significant difference in user allocation because it is primarily

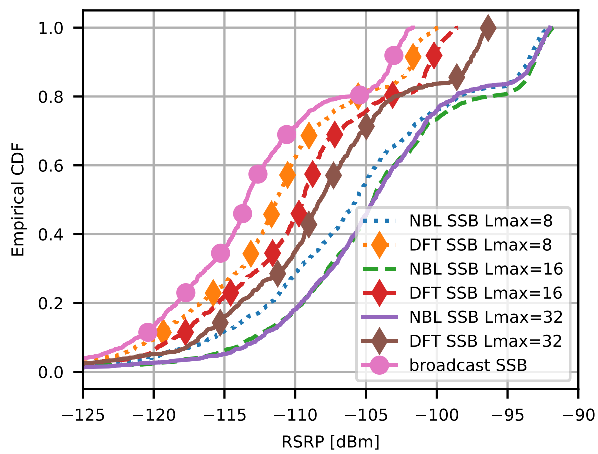

Figure 4.5: The CDF of the reported RSRP using NBL and DFT codebooks of different sizes $L_{\max}$. The broadcast SSB refers to a single, non-beamformed signal and is shown for comparison. The NBL codebooks result in an average of $>$ 5dB gain over DFT codebooks and reduced user outage rates.

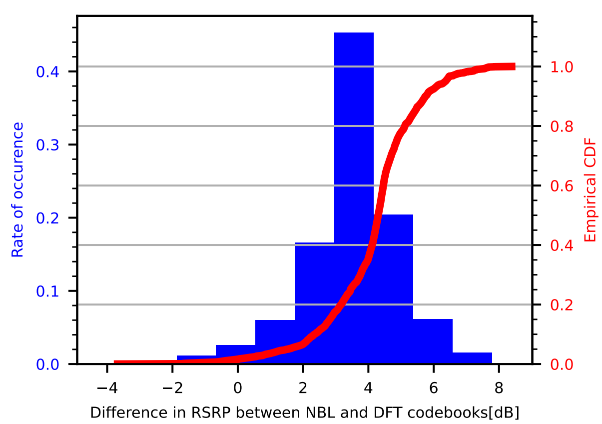

Figure 4.6: A comparison of the reported RSRP by users with NBL codebooks and DFT codebooks. Less than 2% of users reported a decreased RSRP with NBL codebooks, while more than 90% of users gain more than 3dB.

a function of the environment, where Cell 2 experiences NLOS channels for most user locations (see Figure 4.4a).

Table 4.1: User allocation by cell ID

Codebook	Cell 0 User Allocation	Cell 1 User Allocation	Cell 2 User Allocation
NBL	46.6%	45.8%	7.5%
DFT	49.8%	42.2%	8%

4.7.2 CSI-RS performance

While the NBL SSB codebooks noticeably improve the initial access results, evaluating the second goal of supporting CSI-RS with high SE requires a joint evaluation with the CSI-RS. In contrast to initial access, CSI-RS are allocated or designated based on the users already in the network. As such, the candidate CSI-RS codebook should improve the link configuration and support high data rate communications. In the first CSI-RS results, we compare the difference between perfect CSI SU-MIMO SE and the achievable SE from (4.19). The performance difference is equivalent to the training loss function (4.32) and captures the performance and timing relationship as a result of the codebook design process. The results, depicted in Fig. 4.7 highlight the value of AI/ML codebooks over traditional codebooks for beam management, pushing the achievable SE to within 10bps/Hz of the idealized SE in over 80% of results.

While the achievable SE is a critical metric, it is valuable to understand where the SE gains are coming from. The performance gains are especially important if a new system is considered with different algorithms, for example with a different channel estimation technique. Multi-cell interference during CSI-RS causes errors in channel estimation–and subsequent data transmission–so mitigating interference during CSI-RS is important for maximizing performance. We introduce an additional metric, the effective SINR, to help characterize the CSI-RS performance. The effective SINR is determined as the equivalent

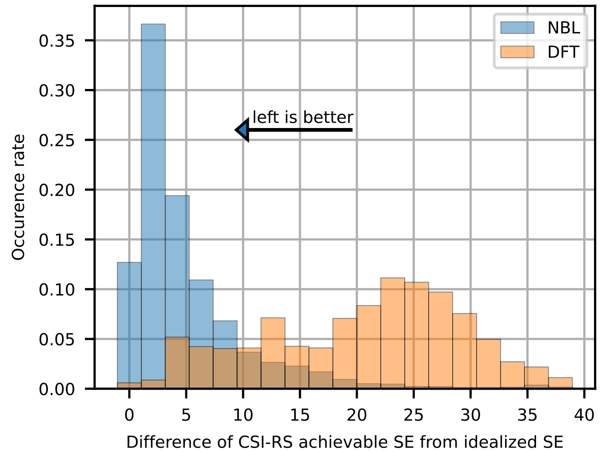

Figure 4.7: A histogram of the difference in achievable and ideal SE. The ideal SE is achieved with perfect CSI and a fully digital array. The difference shows how much SE is lost with respect to the achievable rate during CSI-RS from (4.19). The difference in effective SINR translates to more than 2/3 of users receiving CSI-RS within 5dB of the ideal achievable SE.

SINR needed to obtain the given SE in dB as

$$\text{Eff-SINR}_u = 10 \log_{10}(2^{\text{SE}_u^{\text{CSI-RS}_{\hat{i}_u}}} - 1). \tag{4.33}$$

The effective SINR, which is shown in Figure 4.8, also allows for breaking down the contributions to the achievable rate due to improved signal combining and interference-noise-ratio.

The proposed NBL shows a substantial improvement in effective SINR compared to DFT codebooks. The performance uplift is from the two-fold problem: improving the beamforming coherent combining and reducing the interference from nearby cells during overlapping transmission periods. In fact, if we highlight the contributions between signal strength and inverse interference, the results become even more clear as shown by Figure 4.9. In this case, the effective signal power is significantly improved, although the

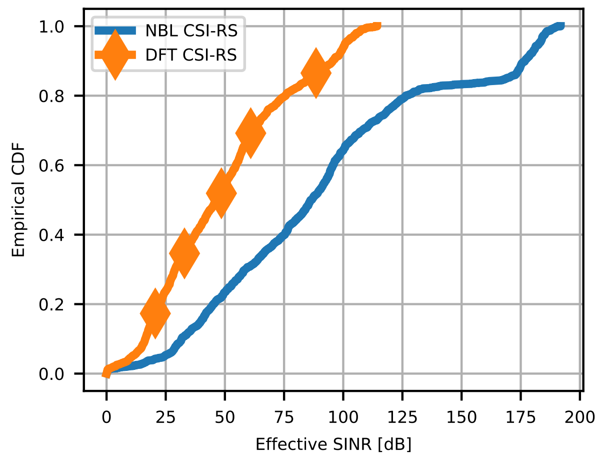

Figure 4.8: A statistical CDF comparison of the effective SINR obtained during CSI-RS reception. It can be seen that the NBL codebooks show a gain of more than 40dB for the majority of users.

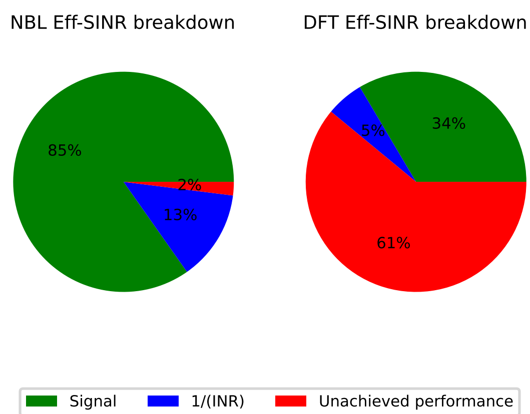

Figure 4.9: A ratio of the average effective signal power and inverse-interference-noise (1/INR) power compared to the achievable effective SINR. The effective SINR difference between NBL and DFT codebooks is significant, but it can be seen that this is a result of both improving signal quality and reducing interference, up to nearly the optimal SINR.

interference is also cut in half.

4.7.3 Network characterization

High achievable SE during CSI-RS, however, does not necessarily imply the wideband, MU-MIMO data transmission will reflect a similar improvement. In the final set of evaluations, we seek to characterize the proposed NBL performance as seen by a network operator. The prior results focused on achievable SE as a result of the beamformed channel–i.e. did not include channel estimation error, quantization error, resources outside of the CSI-RS, or MU-MIMO interference. Although DFT codebooks significantly underperformed during CSI-RS, the beamformed channel with DFT codebooks is more likely to be spatially sparse and see less performance loss from the quantization and limited feedback than learned codebooks. In this subsection, we focus on the ESSE following the entire beam management process. First, the ESSE gain as a result of NBL codebooks in the site-specific case is shown in Figure 4.10. From these results, it is clear there is a valuable performance gain using NBL codebooks, especially considering the ESSE is then multiplied by the channel bandwidth of 100MHz to get the total network rate difference, which is over 2.5Gbps improvement on average. Further, less than 3% of simulations show a loss of performance using the proposed codebooks.

Next, we consider a typical network operator upgrade scenario where data is available from a 5G antenna array (for example $N_{\mathrm{T}} = 64$) that is used to train the model and test with a 6G X-MIMO array ($N_{\mathrm{T}} = 256$) in Figure 4.11. Because of the beamspace transformation, the same neural network can be applied for different arrays which is a valuable benefit for both operators upgrading or potential OpenRAN settings. It can be seen that there is a $\approx$ 10bps/Hz performance loss as a result of the array upscaling, but that the performance is still highly effective compared to generic codebooks.

While array geometry is an interesting generalization, an even more common sce-

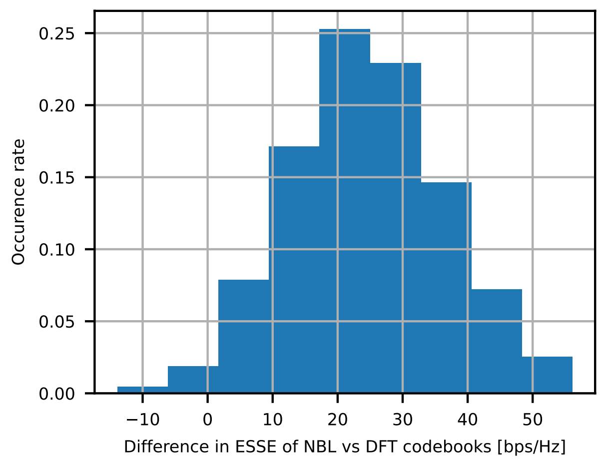

Figure 4.10: Visualization of the difference in ESSE between the proposed NBL and DFT codebooks in the network. Almost universally NBL codebooks show a significant gain in the achievable rate. In comparison, less than 3% of simulations show any decrease in performance.

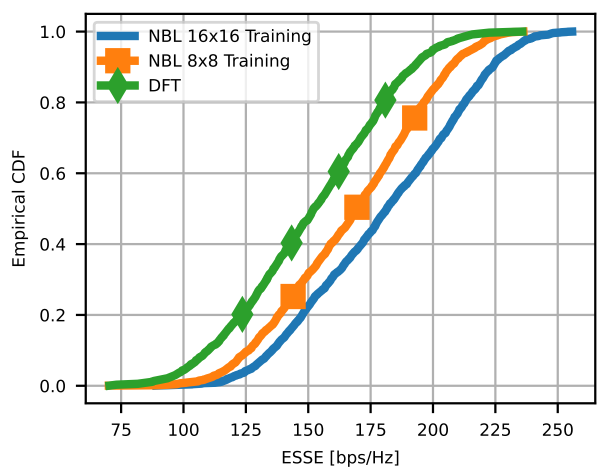

Figure 4.11: Statistically comparing the network ESSE using different training data geometries. In all cases, the codebooks are tested with data corresponding to $N_{\mathrm{X}} = 16, N_{\mathrm{Y}} = 16$ arrays. The benefit of beamspace learning is a natural generalization across geometries allowing for training and testing for different array sizes or even arbitrary configurations (Dreifuerst and Heath 2024a).

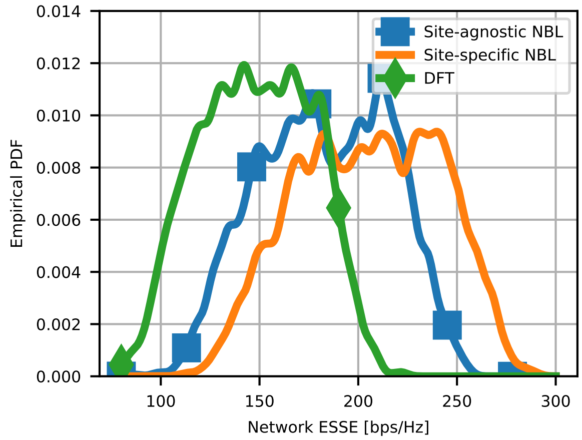

Figure 4.12: The empirical PDF of the network ESSE for NBL and DFT codebooks. The site-agnostic NBL is trained on data from the Munich scene and tested in the Paris scene. The site-specific NBL is trained and tested on data from Paris. It can be seen that NBL generalizes to the new scene to outperform DFT codebooks, although there is still a performance loss.

nario is the generalization across scenes. We propose a comparison where the training and testing raytraced environments are different, and the operating carrier frequency is shifted from 10GHz to 20GHz. Inherently, the beamspace transformation should extend well across frequencies due to the integration of the beamspace transformation with the frequency-dependent array response, however, the frequency change also impacts how the signals interact with the environment. OOD formulations are known to be challenging for AI/ML in wireless (Vuckovic et al. 2023) but even in OOD channels, the network ESSE distribution shows a substantial benefit as a result of NBL codebooks. These results validate both the beamspace representation as a generalizable preprocessing strategy as well as the overall learned model's capabilities.

In the final results, we compare the NBL with prior work (Dreifuerst and Heath 2024a), which was based on learning single-cell power-maximizing codebooks without considering interference. We had previously found that power-maximizing codebooks

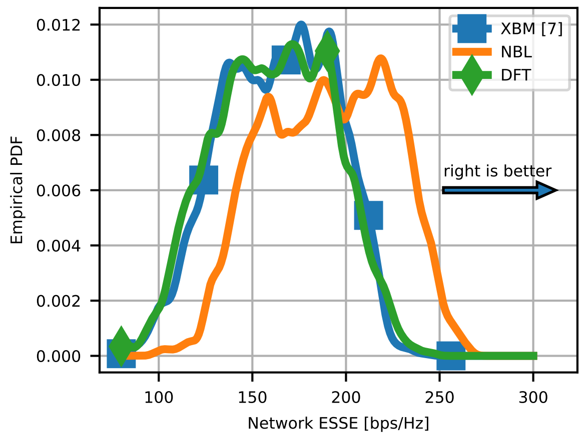

Figure 4.13: A comparison of the XBM model from (Dreifuerst and Heath 2024a) with our proposed solution. The XBM model is given the same training data although each cell operates and learns an independent XBM realization. Our results align with those from (Dreifuerst and Heath 2024a) where the beam alignment gains do not translate to improved ESSE, especially with inter-cell interference.

showed modest gains over DFT codebooks in ESSE, but ultimately the gains were lower than expected based on the beam training performance. We show the relative performance of the two models in Figure 4.13, where the X-BM algorithm from (Dreifuerst and Heath 2024a) is trained independently for each of the three cells. Note that the parameter sizes are not comparable–the prior X-BM algorithm requires over 100 times more parameters. Yet, even with a much smaller network, the change in optimization goal and slight coordination results in a valuable performance improvement. We see that the new NBL formulation overcomes the lack of system performance of the XBM model. All told, the proposed algorithm outperforms traditional codebooks throughout the beam management process and generalizes across a wide range of use cases and deployment strategies.

4.8 Conclusion

In this chapter, we built upon the previous chapters to design a network-wide codebook design algorithm. This algorithm relies on beamspace transformations coupled with a gradient-complete SE metric for end-to-end learning across the network. By introducing interference into the system model and updating the neural network parameters to maximize achievable SE, the learned codebooks are capable of supporting significantly ESSE. The significant network gains also generalize across different array geometries, raytraced environments, and carrier frequencies, showcasing the value of the NBL algorithm over traditional and AI/ML algorithms.

Chapter 5

CONCLUSIONS

This dissertation presents a novel formulation and learning algorithms for codebook design in fully digital, hybrid, and multi-cell hybrid systems. Throughout the chapters, a focus is placed on evaluating the codebook algorithms from a system-level perspective, giving a clear picture of the potential benefits of codebook optimization. In the following sections, we summarize the proposed algorithms and key results of each chapter.

5.1 Chapter 2

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

5.2 Chapter 3

In this chapter, we proposed a machine learning-assisted beam management strategy (XBM) for extreme MIMO. The proposed solution integrates machine learning into all stages of beam management while maintaining realistic timing and computation constraints. ML codebooks provide significant gains during beam training, especially for hybrid arrays, with an average improvement of 6dB in received power compared to traditional methods. In fact, the learned initial access codebooks outperform two-step hierarchical search DFT codebooks in all scenarios.

The proposed XBM algorithm is built upon a beamspace conversion that translates beamformers into a consistent grid representation. We integrate the beamspace conversion, quantization, transmission, and feedback into the model learning to enable an end-to-end learning framework. The proposed algorithm naturally generalizes across array geometries, user distributions, and entirely new environments. Furthermore, XBM does not require changes to UE operation while increasing the achievable spectral efficiency by 10%. We also show that it is also capable of efficiently updating using only the partial CSI available to the network.

Synthesizing the results of our study, we draw the following conclusions about the evolution of the upper-mid band (so-called FR3) in the next releases of 5G and into 6G. First, we find that the beam training framework can be extended to the upper midband and X-MIMO arrays. Second, DFT beams have worse performance as their narrower beamwidths (compared to sub-6GHz bands) become even more suboptimal in multipath environments. Finally, AI/ML approaches are able to design both SSB and CSI-RS codebooks that can substantially improve beam training and sum spectral efficiency within the network. Future investigations on multi-cell and interference-aware codebook design are necessary for further characterizing the benefits of AI/ML in realistic, network-level beam management.

5.3 Chapter 4

In this chapter, we considered a weakly-centralized AI/ML algorithm for beam management in X-MIMO systems. The proposed strategy learns SSB and CSI-RS codebooks for overlapping serving cell regions from beamspace representations. By integrating neural networks into the codebook design process, the cellular network can reduce overhead and obtain stronger wireless links than generic codebooks. Furthermore, codebook design can be integrated with interference-aware metrics to improve overall network performance with hybrid arrays.

The proposed NBL algorithm incorporates an end-to-end learning strategy where codebooks are evaluated and gradients are computed based on the difference with the per-user optimal SE. In combination, the codebook learning relies on beamspace transformations to convert arbitrary beamforming vectors into images representing the magnitude of the array response in a given direction. By integrating this domain knowledge, the NBL can efficiently learn codebooks that generalize across array geometries, environments, and operating bands. We show that the learned codebooks improve network ESSE by more than 25% over traditional codebooks.

The results of this investigation highlight important results with respect to ML/AI for beam management and 6G integration. First, we found that integrating AI/ML into the codebook design process can reduce the need for larger codebooks with X-MIMO arrays. Secondly, we found that although learned beams include arbitrary beamforming directions, the interference can actually be reduced compared to DFT codebooks. This combination of reduced beam training overhead, reduced interference, and improved beam alignment all while generalizing effectively showcases the potential benefits of using AI/ML for beam management.

5.4 Future Work

5.4.1 Deployed model updating

One of the challenges noted in Chapter 3 was the challenge of online updating. In general, the neural networks in this work are trained using drive test data (Lin 2023) with CSI corresponding to the entire physical array. During the results of Figure 4.4, it can be

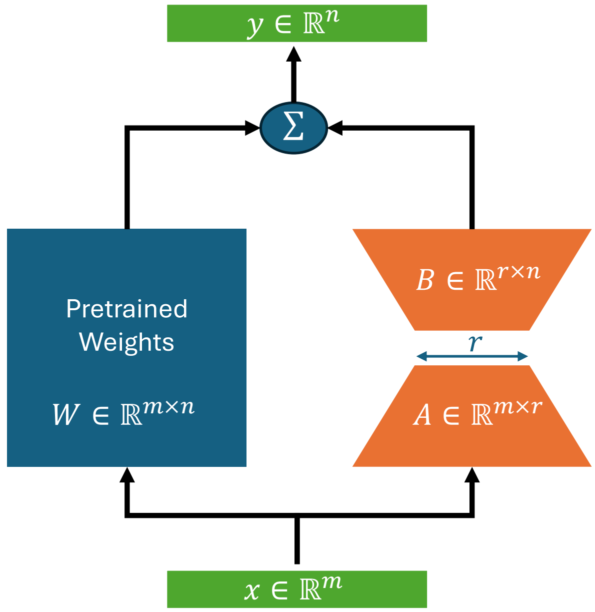

Figure 5.1: LoRA is a low rank model fine-tuning method that trains a secondary model matrix of rank r that modifies the output, rather than the weights, of the pretrained model. While designed for large language models, this formulation also matches the online updates obtained for hybrid beam management models.

seen that only updating using the online information corresponding to the imperfect, low rank digital feedback is suboptimal. One direction of future work would be to investigate a more effective method for model updating from the online data. There has been recent work on using low rank approximations for model updating in large language models (Hu et al. 2022; Liu et al. 2024), but the focus is slightly different as this work is trying to reduce the computational requirements of training as shown in Figure 5.1. Instead, we suggest this strategy may also work when the training data is already a low rank and imperfect approximation.

In addition to low rank updates that can be employed for any neural architecture, specific network structures and training methods have been proposed to address online updating (Rombach et al. 2021; Hu et al. 2022; Liu et al. 2024; Müller et al. 2021; Oala et al. 2022). One useful architecture may be Diffusion models (Rombach et al. 2021) which have been shown to be effective for balancing complexity and precision during retraining. The challenge of updating the model in an online fashion lies in the overlap of the two domains–wireless communications and machine learning–where the wireless network requires very low latency and solutions formulated specifically within the domain yet the ML community has many novel strategies to address model updating (Hu et al. 2022; Liu et al. 2024; Müller et al. 2021; Oala et al. 2022). Therefore online learning solutions will rely on a collaboration of the two domains to achieve an effective and realistic solution.

5.4.2 Heterogeneous hardware beam management

Heterogeneous hardware is an extremely common system model, yet very limited research has developed AI/ML models for wireless communications in these settings, and none have looked at realistic beam management. Heterogeneous hardware introduces algorithmic challenges such as generalization across varying compute budgets, physical layer hardware, and device imperfections. These issues cannot be overlooked as the design of neural networks for wireless requires an understanding of device limitations. Beam management alleviates some of the challenges as the rate of operation is slow relative to typical wireless domains but hardware imperfections or irregular operations by the UEs or even new BS hardware is still a new direction of focus (Peng et al. 2024).

Heterogeneous hardware also introduces systemic challenges as a result of dynamic antenna arrays and geometries. For example, if UEs are equipped with $N_{\mathrm{R}} = \{1, 2, 4, 8, 16\}$ antennas of different geometries, then sparse or ragged vectors are required to represent, concatenate, and allocate information. Although support for ragged and sparse representations is growing in AI/ML, there are still many common functionalities that require additional engineering to enable gradient-based learning with dynamic representations (Lin et al. 2023; Reiser et al. 2021; Fegade et al. 2022). Even in wireless communications, many modern algorithms do not naturally extend to non-uniform or heterogeneous arrays (Xue et al. 2022; Balevi and Andrews 2019; Aghababaiyan et al. 2020). With new array geometries being proposed (Ridwan et al. 2011; Enache et al. 2016; Lin et al. 2020), transformations and preprocessing are necessary for consistent and generalizable learning frameworks with heterogeneous hardware.

5.4.3 Neural CSI quantization

CSI quantization is a standardized process in 5G, although it is notably suboptimal (Wen et al. 2018; Kim et al. 2021; Chen et al. 2022; Xiao et al. 2023). As discussed in Chapter 3 and Chapter 4, CSI quantization has been identified as an important area that AI/ML can improve in 6G networks (Lin 2023). Yet standardizing these methods is difficult as they are site-specific and prone to overfitting on site-specific distributions (Vuckovic et al. 2023; Wen et al. 2018). This level of overfitting is not necessarily problematic so long as distributions are stationary or slowly changing. The challenge is that any encoder-decoder models that are used by the UE require significant computation to meet the CSI reporting periods (Giordani et al. 2019; Dreifuerst and Heath 2023b), while also experiencing distribution drift (Oala et al. 2022) as a result of mobility throughout a network. The sharing, training, and resharing of models is an inherent challenge of neural networks for CSI quantization.

One-sided CSI models have also been proposed (Lee et al. 2023; Chen et al. 2023), where the CSI quantization is standardized but the BS translates the feedback prior to reconstruction (2.20). The typical goal of these models is to reduce the influence of noise, which is often significant in the uplink transmission due to link budget asymmetry

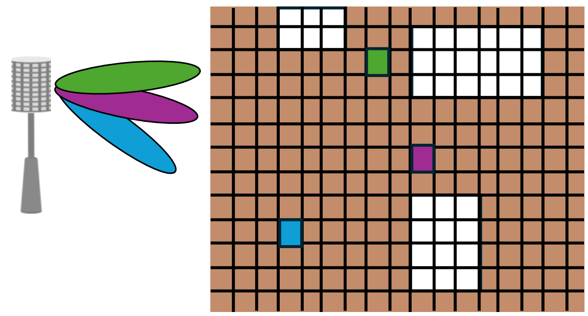

Figure 5.2: An ISAC system using sensing information to identify and predict beamformers. Users are identified by colored locations in the occupancy grid while buildings are shown by white grids. ISAC allows for efficient resource usage while obtaining valuable environmental context for beam management.

(Becirovic et al. 2022). Yet the value of SNR boosting during CSI feedback reception is less critical in X-MIMO regimes due to the increased array sizes. Alternatively, the models can be designed to improve the multipath channel information rather than simply denoising. Such end-to-end formulations are dependent on the entire beam management process, however, making it a challenging direction of research in the face of distribution shift, online updating, and heterogenous hardware as mentioned previously. Integrating the complete beam management process into an end-to-end framework with CSI quantization has the potential to utilize AI/ML to achieve even greater gains than those shown in Chapter 4.

5.4.4 ISAC-assisted beam management

Finally, integrated sensing and communications is also a relevant direction of work identified for 6G (Lin 2023) and within academia (Dias et al. 2019; Zheng and González-Prelcic 2019; Ali et al. 2021; Rezaie et al. 2022). ISAC focuses on reusing the wireless signals to also measure positioning and environmental reflections. Straightforward ISAC beam training strategies rely on positioning from ISAC systems to improve beamforming (Va et al. 2017; Dias et al. 2019; Cui et al. 2019). Other ideas for gathering and integrating environmental awareness have been proposed (Y. Wang et al. 2019; Nishio et al. 2021). The idea of using sensing to obtain occupancy grids, location-specific modeling, and environmental context presents new possibilities for mobility robustness and low-latency applications. For example, an occupancy grid identifying users and environment allows for predicting beamformers for users with low overhead as shown in Figure 5.2. In fact, ISAC can present an interesting solution to the neural CSI quantization challenges by incorporating spatial models to implicitly capture directional CSI. Therefore beam training overhead and even feedback overhead can be reduced while improving network ESSE.

REFERENCES

3GPP (2023a). TS 38.214 V16.7.1 NR; Physical layer procedures for data.

3GPP (2023b). WI summary for WI Core part: NR MIMO evolution for downlink and uplink.

Aghababaiyan, K., Shah-Mansouri, V., and Maham, B. (2020). High-precision omp-based direction of arrival estimation scheme for hybrid non-uniform array. *IEEE Commun. Lett.*, 24(2):354–357.

Ali, A., Mo, J., Ng, B. L., Va, V., and Zhang, J. C. (2021). Orientation-Assisted Beam Management for beyond 5G Systems. *IEEE Access*, 9:51832–51846.

Alkhateeb, A., El Ayach, O., Leus, G., and Heath, R. W. (2014). Channel estimation and hybrid precoding for millimeter wave cellular systems. *IEEE J. Sel. Top. Signal Process.*, 8(5):831–846.

Alouzi, M., Al-Kamali, F., D'amours, C., and Chan, F. (2023). Direct conversion of hybrid precoding and combining from full array architecture to subarray architecture for mmWave MIMO systems. *IEEE Access*, 11:35457–35468.

Balevi, E. and Andrews, J. G. (2019). Online antenna tuning in heterogeneous cellular networks with deep reinforcement learning. *IEEE Transactions on Cognitive Communications and Networking*, 5(4):1113–1124.

Becirovic, E., Bjornson, E., and Larsson, E. G. (2022). Combining reciprocity and CSI feedback in MIMO systems. *IEEE Trans. Wireless Commun.*

Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. (2010). A theory of learning from different domains. *Machine Learning*, 79:151–175.

Castellanos, M. R., Raghavan, V., Ryu, J. H., Koymen, O. H., Li, J., Love, D. J., and Peleato, B. (2018). Channel-Reconstruction-Based Hybrid Precoding for Millimeter-Wave Multi-User MIMO Systems. *IEEE J. Sel. Areas Commun.*, 12(2):383–398.

Chen, M., Guo, J., Wen, C. K., Jin, S., Li, G. Y., and Yang, A. (2022). Deep learning-based implicit CSI feedback in massive MIMO. *IEEE Trans. Commun.*, 70:935–950.

Chen, W., Wan, W., Wang, S., Sun, P., Li, G. Y., and Ai, B. (2023). CSI-PPPNet: A One-Sided One-for-All Deep Learning Framework for Massive MIMO CSI Feedback. *IEEE Trans. Wireless Commun.*, pages 1–1.

Choi, J., Va, V., Gonzalez-Prelcic, N., Daniels, R., Bhat, C. R., and Heath, R. W. (2016). Millimeter-wave vehicular communication to support massive automotive sensing. *IEEE Communications Magazine*, 54(12):160–167.

Cui, Y., Xu, H., Wu, J., Sun, Y., and Zhao, J. (2019). Automatic vehicle tracking with roadside LiDAR data for the connected-vehicles system. *IEEE Intelligent Systems*, 34(3):44–51.

Dias, M., Klautau, A., González-Prelcic, N., and Heath, R. W. (2019). Position and LIDAR-Aided mmWave Beam Selection using Deep Learning. In *2019 IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 1–5.

Dreifuerst, R., Daulton, S., Qian, Y., Varkey, P., Balandat, M., Sanjay, K., Anoop, T., Ali, Y., Vish, P., and Robert W., H. (2021). Optimizing coverage and capacity in cellular networks using machine learning. In *Proc of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.

Dreifuerst, R. M. and Heath, R. W. (2022). SignalNet: A low resolution sinusoid decomposition and estimation network. *IEEE Trans. Signal Process.*, 70:4454–4467.

Dreifuerst, R. M. and Heath, R. W. (2023a). Machine learning codebook design for initial access and CSI type-II feedback in sub-6GHz 5G NR. *IEEE Trans. Wireless Commun.*

Dreifuerst, R. M. and Heath, R. W. (2023b). Massive MIMO in 5G: How beamforming, codebooks, and feedback enable larger arrays. *IEEE Commun. Mag.*

Dreifuerst, R. M. and Heath, R. W. (2024a). Hierarchical ML codebook design for extreme MIMO beam management. *Submitted to IEEE Transactions on Machine Learning in Communications and Networking.*

Dreifuerst, R. M. and Heath, R. W. (2024b). Neural codebook design for network beam management. *Submitted to IEEE Trans. Wireless Commun.*

Dreifuerst, R. M., Heath, R. W., Kulkarni, M. N., and Charlie, J. (2020a). Deep Learning-based Carrier Frequency Offset Estimation with One-Bit ADCs. In *2020 IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 1–5.

Dreifuerst, R. M., Heath, R. W., Kulkarni, M. N., and Charlie, J. (2020b). Deep learning-based carrier frequency offset estimation with one-bit ADCs. In *IEEE Work. Signal Process. Adv. Wirel. Commun. (SPAWC)*.

Dreifuerst, R. M., Heath, R. W., and Yazdan, A. (2022). Massive MIMO beam management in sub-6 GHz 5G NR. In *Proc. of the IEEE Veh. Technol. Conf. (VTC)*, pages 1–5.

Dreifuerst, R. M. and Heath, R. W. J. (2023c). CSI type-II codebook of codebooks. In *Proc. of the IEEE Work. Signal Process. Adv. Wirel. Commun. (SPAWC)*.

Dreifuerst, R. M., Kilinc, I., and Heath, R. W. (2024). Context-aware codebook design for 6g extreme mimo. In preparation for submission to IEEE SPAWC.

Enache, F., Popescu, F., Depărăţeanu, D., Enache, A., and Oroian, T. (2016). Multi-criteria optimization of non-uniform linear antenna array using genetic algorithms. In *International Conference on Communications (COMM)*, pages 121–124.

Fegade, P., Chen, T., Gibbons, P. B., and Mowry, T. C. (2022). The CoRa tensor compiler: Compilation for ragged tensors with minimal padding.

Felix, A., Cammerer, S., Dörner, S., Hoydis, J., and Ten Brink, S. (2018). Ofdm-autoencoder for end-to-end learning of communications systems. In *Proc. of the IEEE Work. Signal Process. Adv. Wirel. Commun. (SPAWC)*, pages 1–5.

Giordani, M., Polese, M., Roy, A., Castor, D., and Zorzi, M. (2019). A tutorial on beam management for 3GPP NR at mmwave frequencies. *IEEE Commun. Surv. Tutorials*, 21(1):173–196.

Godara, L. (1997). Application of antenna arrays to mobile communications. ii. beamforming and direction-of-arrival considerations. *Proceedings of the IEEE*, 85(8):1195–1245.

Goutay, M., Aoudia, F. A., Hoydis, J., and Gorce, J.-M. (2020). Machine Learning for MU-MIMO Receive Processing in OFDM Systems.

Gupta, M., Dreifuerst, R. M., Yazdan, A., Huang, P.-H., Kasturia, S., and Andrews, J. G. (2021). Load balancing and handover optimization in multi-band networks using deep reinforcement learning. In *IEEE Global Communications Conference (GLOBECOM)*, pages 1–6.

Hamid, S., Chopra, S. R., Gupta, A., Tanwar, S., Florea, B. C., Taralunga, D. D., Alfarraj, O., and Shehata, A. M. (2023). Hybrid Beamforming in Massive MIMO for Next-Generation Communication Technology. *Sensors*, 23(16):7294.

Heath, R. W., Gonzalez-Prelcic, N., Rangan, S., Roh, W., and Sayeed, A. M. (2016). An overview of signal processing techniques for millimeter wave MIMO systems. *IEEE J. Sel. Top. Signal Process.*, 10(3):436–453.

Heath Jr., R. W. and Lozano, A. (2018). *Foundations of MIMO Communication*. Cambridge University Press.

Heng, Y., Andrews, J. G., Mo, J., Va, V., Ali, A., Ng, B. L., and Zhang, J. C. (2021). Six key challenges for beam management in 5.5G and 6G systems. *IEEE Commun. Mag.*, 59(7):74–79.

Heng, Y., Mo, J., and Andrews, J. G. (2022). Learning site-specific probing beams for fast mmWave beam alignment. *IEEE Trans. Wirel. Commun.*, 21(8):5785–5800.

Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models.

Holma, H., Viswanathan, H., and Mogensen, P. (2021). Extreme massive MIMO for macro cell capacity boost in 5G-Advanced and 6G. Technical report, Nokia Bell Labs.

Hoydis, J., Cammerer, S., Ait Aoudia, F., Vem, A., Binder, N., Marcus, G., and Keller, A. (2022). Sionna: An open-source library for next-generation physical layer research. *arXiv preprint*.

Hu, E., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. (2022). LORA: Low-rank adaptation of large language models. In *ICLR 2022 - 10th Int. Conf. Learn. Represent.* International Conference on Learning Representations, ICLR.

Hussain, M. and Michelusi, N. (2022). Learning and adaptation for millimeter-wave beam tracking and training: A dual timescale variational framework. *IEEE J. Sel. Areas Commun.*, 40(1):37–53.

Jaeckel, S., Raschkowski, L., Börner, K., and Thiele, L. (2014). QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials. *IEEE Trans. Antennas Propag.*, pages 3242–3256.

Jiang, C., Zhang, H., Ren, Y., Han, Z., Chen, K. C., and Hanzo, L. (2017). Machine Learning Paradigms for Next-Generation Wireless Networks. *IEEE Wirel. Commun.*, 24(2):98–105.

Kilinc, I., Dreifuerst, R. M., Kim, J., and Heath, R. W. (2024). Beam alignment in mmWave vehicular systems: Machine learning for beam decoupling. In preparation for submission to IEEE SPAWC.

Kim, J., Lee, H., and Park, S.-H. (2021). Learning robust beamforming for MISO downlink systems. *IEEE Communications Letters*, 25(6):1916–1920.

Kingma, D. and Ba, J. (2014). Adam: A method for stochastic optimization. *International Conference on Learning Representations*.

Lee, A., Park, H., Kwon, Y., Lee, H., and Chong, S. (2023). Practical Denoising Autoencoder for CSI Feedback Without Clean Target in Massive MIMO Networks. *IEEE Wirel. Commun. Lett.*

Lee, K., Kim, J., Jin, E. W., and Kim, K. S. (2022). Extreme massive mimo for upper-mid band 6g communications. In *2022 13th International Conference on Information and Communication Technology Convergence (ICTC)*, pages 997–999.

Li, B., Zhou, Z., Zou, W., Sun, X., and Du, G. (2013). On the efficient beam-forming training for 60GHz wireless personal area networks. *IEEE Trans. Wireless Commun.*, 12(2):504–515.

Li, K., Li, Y., Cheng, L., Shi, Q., and Luo, Z. Q. (2021). Learning enhanced beamforming vector from CQIs in 5G NR FDD massive MIMO systems: A tuning-free approach. In *Proc. of the IEEE Work. Signal Process. Adv. Wirel. Commun. (SPAWC)*.

Li, L., Zeng, X., Liu, Y. F., Xu, Y., and Chang, T. H. (2023). CSI sensing from heterogeneous user feedbacks: A constrained phase retrieval approach. *IEEE Trans. Wireless Commun.*

Li, X., Jin, S., Gao, X., and Heath, R. W. (2016). Three-dimensional beamforming for large-scale FD-MIMO systems exploiting statistical channel state information. *IEEE Trans. Veh. Technol.*, 65(11):8992–9005.

Lin, X. (2023). An overview of the 3GPP study on artificial intelligence for 5G New Radio.

Lin, Y.-C., Lee, T.-S., and Ding, Z. (2023). A Scalable Deep Learning Framework for Dynamic CSI Feedback with Variable Antenna Port Numbers. *IEEE Trans. Wireless Commun.*, pages 1–1.

Lin, Z., Lv, T., Ni, W., Zhang, J. A., and Liu, R. P. (2020). Tensor-based multi-dimensional wideband channel estimation for mmwave hybrid cylindrical arrays. *IEEE Trans. Commun.*, 68(12):7608–7622.

Liu, S.-Y., Wang, C.-Y., Yin, H., Molchanov, P., Wang, Y.-C. F., Cheng, K.-T., and Chen, M.-H. (2024). DoRA: Weight-decomposed low-rank adaptation.

Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. of the IEEE Int. Conf. Comput. Vis.*, pages 9992–10002.

Lopes, V. H. L., Nahum, C. V., Dreifuerst, R. M., Batista, P., Klautau, A., Cardoso, K. V., and Heath, R. W. (2022). Deep reinforcement learning-based scheduling for multiband massive MIMO. *IEEE Access*, 10:125509–125525.

Ma, K., Sang, Y., Ming, Y., Lian, J., Tian, C., and Wang, Z. (2023). Improving the Performance of R17 Type-II Codebook with Deep Learning.

MathWorks (2023). *5G Toolbox*.

Morozov, G., Davydov, A., and Sergeev, V. (2016). Enhanced CSI feedback for FD-MIMO with beamformed CSI-RS in LTE-A Pro systems. In *Proc. of the IEEE Veh. Technol. Conf.*, pages 1–5.

Müller, T., Fabrice Rousselle, N., Jan Novák, N., Alexander Keller, N., Rousselle, F., Novák, J., and Keller, A. (2021). Real-time neural radiance caching for path tracing. *ACM Trans. Graph*, 40(4):36.

Nahum, C. V., Lopes, V. H. L., Dreifuerst, R. M., Batista, P., Correa, I., Cardoso, K. V., Klautau, A., and Heath, R. W. (2024). Intent-aware radio resource scheduling in a RAN slicing scenario using reinforcement learning. *IEEE Trans. Wireless Commun.*, 23(3):2253–2267.

Nishio, T., Koda, Y., Park, J., Bennis, M., and Doppler, K. (2021). When wireless communications meet computer vision in beyond 5G. *IEEE Communications Standards Magazine*, 5(2):76–83.

Oala, L., Aversa, M., Willis, K., Nobis, G., Neuenschwander, Y., Buck, M., Matek, C., Munich, H. Z., Extermann, J., Pomarico, E., Samek, W., Murray-Smith, R., Clausen, C., and Sanguinetti, B. (2022). Data models for dataset drift controls in machine learning with images.

Peng, Z., Zhang, Z., Pan, C., Renzo, M. D., Dobre, O. A., and Wang, J. (2024). Beamforming optimization for active RIS-aided multiuser communications with hardware impairments.

Qin, Z. and Yin, H. (2023). A review of codebooks for CSI feedback in 5G New Radio and beyond.

Raghavan, V., Cezanne, J., Subramanian, S., Sampath, A., and Koymen, O. (2016). Beamforming Tradeoffs for Initial UE Discovery in Millimeter-Wave MIMO Systems. *IEEE J. Sel. Top. Signal Process.*, 10(3):543–559.

Rebato, M., Polese, M., and Zorzi, M. (2018). Multi-sector and multi-panel performance in 5G mmWave cellular networks. In *IEEE GLOBECOM*, pages 1–6.

Reddi, V. J., Cheng, C., Kanter, D., Mattson, P., Schmuelling, G., Wu, C.-J., Anderson, B., Breughe, M., Charlebois, M., Chou, W., Chukka, R., Coleman, C., Davis, S., Deng, P., Diamos, G., Duke, J., Fick, D., Gardner, J. S., Hubara, I., ..., and Zhou, Y. (2020). MLPerf inference benchmark. V3.1 updated Aug. 2023.

Reiser, P., Eberhard, A., and Friederich, P. (2021). Graph neural networks in TensorFlow-Keras with raggedtensor representation (kgcnn). *Software Impacts*, 9:100095.

Ren, Y., Yoo, S., and Hoisie, A. (2019). Performance analysis of deep learning workloads on leading-edge systems. In *2019 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pages 103–113.

Rezaie, S., Morais, J., de Carvalho, E., Alkhateeb, A., and Manchón, C. N. (2022). Location- and orientation-aware millimeter wave beam selection for multi-panel antenna devices. In *IEEE Global Communications Conference*, pages 597–602.

Ridwan, M., Abdo, M., and Jorswieck, E. (2011). Design of non-uniform antenna arrays using genetic algorithm. In *International Conference on Advanced Communication Technology (ICACT)*, pages 422–427.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2021). High-Resolution Image Synthesis with Latent Diffusion Models. *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695.

Samsung (2020). Massive MIMO for new radio.

Sayeed, A. (2002). Deconstructing multiantenna fading channels. *IEEE Trans. Signal Process.*, 50(10):2563–2579.

Shafin, R., Shafin, R., Chen, H., Nam, Y. H., Nam, Y. H., Hur, S., Park, J., Zhang, J., Reed, J. H., and Liu, L. (2020). Self-tuning sectorization: Deep reinforcement learning meets broadcast beam optimization. *IEEE Trans. Wireless Commun.*, 19(6):4038–4053.

Sim, G. H., Klos, S., Asadi, A., Klein, A., and Hollick, M. (2018). An online context-aware machine learning algorithm for 5g mmwave vehicular communications. *IEEE/ACM Transactions on Networking*, 26(6):2487–2500.

Singh, A. and Joshi, S. (2021). A survey on hybrid beamforming in mmWave massive MIMO system. *J. Sci. Res.*, 65(01):201–213.

Turan, N., Fesl, B., Joham, M., Ma, Z., Soong, A. C. K., Sheen, B., Xiao, W., and Utschick, W. (2024). Limited Feedback on Measurements: Sharing a Codebook or a Generative Model?

V. Va, J. Choi, and Heath, R. W. (2017). The Impact of Beamwidth on Temporal Channel Variation in Vehicular Channels and Its Implications. *IEEE Transactions on Vehicular Technology*, 66(6):5014–5029.

V. Va, J. Choi, T. Shimizu, G. Bansal, and Heath, R. W. (2018). Inverse multipath fingerprinting for millimeter wave V2I beam alignment. *IEEE Trans. Veh. Technol.*, 67(5):4042–4058.

Va, V., Shimizu, T., Bansal, G., and Heath, R. W. (2017). Position-aided millimeter wave V2I beam alignment: A learning-to-rank approach. In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 1–5.

Va, V., Vikalo, H., and Heath, R. W. (2016). Beam tracking for mobile millimeter wave communication systems. In *2016 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 743–747.

Vuckovic, K., Mashhadi, M. B., Hejazi, F., Rahnavard, N., and Alkhateeb, A. (2023). PARAMOUNT: Towards generalizable deeP leARning for mmwAve beaM selectiOn using sUb-6GHz chaNnel measuremenTs. *IEEE Trans. Wireless Commun.*

Wang, J., Lan, Z., woo Pyo, C., Baykas, T., Sum, C., Rahman, M., Gao, J., Funada, R., Kojima, F., Harada, H., and Kato, S. (2009). Beam codebook based beamforming protocol for multi-Gbps millimeter-wave WPAN systems. *IEEE J. Sel. Areas Commun.*, 27(8):1390–1399.

Wang, T., Wen, C. K., Wang, H., Gao, F., Jiang, T., and Jin, S. (2017). Deep learning for wireless physical layer: Opportunities and challenges. *China Commun.*, 14(11):92–111.

Wang, Y., Myers, N. J., González-Prelcic, N., and Heath, R. W. (2021). Site-specific online compressive beam codebook learning in mmWave vehicular communication. *IEEE Trans. Wireless Commun.*, pages 1–14.

Wen, C. K., Shih, W. T., and Jin, S. (2018). Deep learning for massive MIMO CSI feedback. *IEEE Wirel. Commun. Lett.*, 7(5):748–751.

Wu, X., Liu, D., and Yin, F. (2018). Hybrid beamforming for multi-user massive MIMO systems. *IEEE Trans. Commun.*, 66(9):3879–3891.

Xia, W., Zheng, G., Zhu, Y., Zhang, J., Wang, J., and Petropulu, A. P. (2020). A deep learning framework for optimization of MISO downlink beamforming. *IEEE Trans. Commun.*, 68(3):1866–1880.

Xiao, H., Tian, W., Liu, W., Guo, J., Zhang, Z., Jin, S., Shi, Z., Guo, L., and Shen, J. (2023). Knowledge-driven meta-learning for CSI feedback.

Xiao, Z., He, T., Xia, P., and Xia, X.-G. (2016). Hierarchical codebook design for beamforming training in millimeter-wave communication. *IEEE Trans. Wireless Commun.*, 15(5):3380–3392.

Xue, Q., Ji, C., Ma, S., Guo, J., Xu, Y., Chen, Q., and Zhang, W. (2023). A survey of beam management for mmWave and THz communications towards 6G.

Xue, Q., Liu, Y. J., Sun, Y., Wang, J., Yan, L., Feng, G., and Ma, S. (2022). Beam management in ultra-dense mmWave network via federated reinforcement learning: An intelligent and secure approach. *IEEE Trans. Cogn. Commun. Netw.*

Y. Wang, A. Klautau, M. Ribero, A. C. K. Soong, and Heath, R. W. (2019). MmWave vehicular beam selection with situational awareness using machine learning. *IEEE Access*, 7:87479–87493.

Yang, J., Ge, X., and Zhong, Y. (2019). How much of wireless rates can smartphones support in 5G networks? *IEEE Network*, 33(3):122–129.

Yang, J., Zhu, W., and Tao, M. (2022). Deep learning for hierarchical beam alignment in mmWave communication systems. In *Proc. of the IEEE Global Communications Conference*.

Yang, J., Zhu, W., Tao, M., and Sun, S. (2023). Hierarchical Beam Alignment for Millimeter-Wave Communication Systems: A Deep Learning Approach. *IEEE Trans. Wireless Commun.*, pages 1–1.

Zhang, C., Patras, P., and Haddadi, H. (2019). Deep Learning in Mobile and Wireless Networking: A Survey. *IEEE Commun. Surv. Tutorials*, 21(3):2224–2287.

Zhang, Y. and Alkhateeb, A. (2024). Decentralized Interference-Aware Codebook Learning in Millimeter Wave MIMO Systems.

Zheng, W. and González-Prelcic, N. (2019). Joint position, orientation AND channel estimation in hybrid mmWAVE MIMO systems. In *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pages 1453–1458.

Zhou, P., Fang, X., Wang, X., Long, Y., He, R., and Han, X. (2019). Deep learning-based beam management and interference coordination in dense mmWave networks. *IEEE Trans. Veh. Technol.*, 68(1):592–603.

APPENDIX

Appendix A

Publications

A.1 Publications related to the dissertation

1. Dreifuerst, R. M. and Heath, R. W. (2023b). Massive MIMO in 5G: How beamforming, codebooks, and feedback enable larger arrays. *IEEE Commun. Mag.*

2. Dreifuerst, R. M. and Heath, R. W. J. (2023c). CSI type-II codebook of codebooks. In *Proc. of the IEEE Work. Signal Process. Adv. Wirel. Commun. (SPAWC).*

3. Dreifuerst, R. M. and Heath, R. W. (2023a). Machine learning codebook design for initial access and CSI type-II feedback in sub-6GHz 5G NR. *IEEE Trans. Wireless Commun.*

4. Dreifuerst, R. M. and Heath, R. W. (2024a). Hierarchical ML codebook design for extreme MIMO beam management. *Submitted to IEEE Transactions on Machine Learning in Communications and Networking.*

5. Dreifuerst, R. M., Heath, R. W., and Yazdan, A. (2022). Massive MIMO beam management in sub-6 GHz 5G NR. In *Proc. of the IEEE Veh. Technol. Conf.*

(VTC), pages 1–5.

6. Dreifuerst, R. M. and Heath, R. W. (2024b). Neural codebook design for network beam management. *Submitted to IEEE Trans. Wireless Commun.*.

7. Dreifuerst, R. M., Kilinc, I., and Heath, R. W. (2024). Context-aware codebook design for 6g extreme mimo. In preparation for submission to IEEE SPAWC.

A.2 Additional publications

1. Dreifuerst, R., Daulton, S., Qian, Y., Varkey, P., Balandat, M., Sanjay, K., Anoop, T., Ali, Y., Vish, P., and Robert W., H. (2021). Optimizing coverage and capacity in cellular networks using machine learning. In *Proc of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.

2. Dreifuerst, R. M., Heath, R. W., Kulkarni, M. N., and Charlie, J. (2020a). Deep Learning-based Carrier Frequency Offset Estimation with One-Bit ADCs. In *2020 IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 1–5.

3. Dreifuerst, R. M. and Heath, R. W. (2022). SignalNet: A low resolution sinusoid decomposition and estimation network. *IEEE Trans. Signal Process.*, 70:4454–4467.

4. Dreifuerst, R. M., Heath, R. W., Kulkarni, M. N., and Charlie, J. (2020b). Deep learning-based carrier frequency offset estimation with one-bit ADCs. In *IEEE Work. Signal Process. Adv. Wirel. Commun. (SPAWC)*.

5. Kilinc, I., Dreifuerst, R. M., Kim, J., and Heath, R. W. (2024). Beam alignment in mmWave vehicular systems: Machine learning for beam decoupling. In preparation

for submission to IEEE SPAWC.

6. Nahum, C. V., Lopes, V. H. L., Dreifuerst, R. M., Batista, P., Correa, I., Cardoso, K. V., Klautau, A., and Heath, R. W. (2024). Intent-aware radio resource scheduling in a RAN slicing scenario using reinforcement learning. *IEEE Trans. Wireless Commun.*, 23(3):2253–2267.

7. Lopes, V. H. L., Nahum, C. V., Dreifuerst, R. M., Batista, P., Klautau, A., Cardoso, K. V., and Heath, R. W. (2022). Deep reinforcement learning-based scheduling for multiband massive MIMO. *IEEE Access*, 10:125509–125525.

8. Gupta, M., Dreifuerst, R. M., Yazdan, A., Huang, P.-H., Kasturia, S., and Andrews, J. G. (2021). Load balancing and handover optimization in multi-band networks using deep reinforcement learning. In *IEEE Global Communications Conference (GLOBECOM)*, pages 1–6.