

Adaptive Neural Signal Detection for Massive MIMO

by

Mehrdad Khani Shirkoohi

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Signature redacted

Author

Department of Electrical Engineering and Computer Science

May 24, 2019

Signature redacted

Certified by

Mohammad Alizadeh

Assistant Professor of Electrical Engineering and Computer Science

Thesis Supervisor

Signature redacted

Accepted by

Leslie A. Kolodziejski

Professor of Electrical Engineering and Computer Science

Chair, Department Committee on Graduate Students



ARCHIVES

Adaptive Neural Signal Detection for Massive MIMO

by

Mehrdad Khani Shirkoohi

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2019, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

Massive Multiple-Input Multiple-Output (MIMO) is a key enabler for fifth generation (5G) cellular communication systems. Massive MIMO gives rise to challenging signal detection problems for which traditional detectors are either impractical or suffer from performance limitations. Recent work has proposed several learning approaches to MIMO detection with promising results on simple channel models (e.g., i.i.d. Gaussian entries). However, we find that the performance of these schemes degrades significantly in real-world scenarios in which the channels of different receivers are spatially correlated. The root of this poor performance is that these schemes either do not exploit the problem structure (requiring models with millions of training parameters), or are overly-constrained to mimic algorithms that require very specific assumptions about the channel matrix.

We propose MMNet, a deep learning MIMO detection scheme that significantly outperforms existing approaches on realistic channel matrices with the same or lower computational complexity. MMNet's design builds on the theory of iterative soft-thresholding algorithms to identify the right degree of model complexity, and it uses a novel training algorithm that leverages temporal and frequency locality of channel matrices at a receiver to accelerate training. Together, these innovations allow MMNet to train online for every realization of the channel. On i.i.d. Gaussian channels, MMNet requires 2 orders of magnitude fewer operations than existing deep learning schemes but achieves near-optimal performance. On spatially-correlated realistic channels, MMNet achieves the same error rate as the next-best learning scheme (OAMPNet [1]) at 2.5dB lower Signal-to-Noise Ratio (SNR) and with at least $10\times$ less computational complexity. MMNet is also 4-8dB better overall than a classic linear scheme like the minimum mean square error (MMSE) detector.

Thesis Supervisor: Mohammad Alizadeh

Title: Assistant Professor of Electrical Engineering and Computer Science

Acknowledgments

I thank my advisor, Mohammad, for being a true patient friend. It has been a blessing for me to work with him during the past years, and he has always been a great source of inspiration for me.

I would also like to thank my academic advisor, Guy Bresler, for his thoughtful and flexible support.

This thesis is the result of joint work with Mohammad Alizadeh and our collaborators from Nokia Bell Labs, Jakob Hoydis, and Phil Fleming. I learned a lot from my collaborators during this project, and I am thankful for that.

I have been lucky to have great labmate and friend, Peter Iannucci, for his generous sharing of valuable experiences and thoughtful suggestions. I also like to thank our lab members and friends Hongzi Mao, Prateesh Goyal, Ravichandra Addanki, Songtao He, Amy Ousterhout, Akshay Narayan, Frank Cangialosi, Vikram Nathan, Venkat Arun, Vibhaalakshmi Sivaraman, Shaileshh Bojja Venkatakrishnan, and Sheila Marian.

I am thankful to MIT and specially CSAIL for providing such an excellent environment for me to pursue my passions.

Lastly, but most importantly, I am thankful to my family. I am grateful to my parents and sister for their unconditional love and support. My special thanks go to Zahra, for all her love and support.

Contents

1	Introduction	13
2	Background and Related Work	17
2.1	Problem Definition	17
2.2	An iterative framework for MIMO detection	18
2.3	Classical MIMO detection algorithms	20
2.3.1	Linear	20
2.3.2	Approximate Message Passing (AMP)	21
2.3.3	Other techniques	22
2.4	Learning-based MIMO detection schemes	23
2.4.1	DetNet	23
2.4.2	OAMPNet	24
3	MMNet Design	25
4	Experiments	29
4.1	Methodology	30
4.1.1	Compared Schemes	30
4.1.2	Dataset	31
4.1.3	Training	32
4.2	Results	33
4.2.1	independent and identically distributed (i.i.d.) Gaussian channels	33
4.2.2	Realistic channels	35

5	Why MMNet works	39
5.1	Error dynamics	39
5.2	Analysis	40
5.3	Impact of channel condition number	44
6	Online Training Algorithm	47
6.1	Channel locality	48
6.2	Training algorithm and results	49
6.3	Computational complexity	51
7	Conclusion	53

List of Figures

- 2-1 A block of an iterative detector in general framework. Each block contains a linear transformation followed by a denoising stage. 18

- 4-1 symbol error rate (SER) vs. signal-to-noise ratio (SNR) of different schemes for three modulations (QAM4, QAM16 and QAM64) and two system sizes (32 and 16 transmitters, 64 receivers) with i.i.d. Gaussian channels. 34
- 4-2 SER vs. SNR of different schemes for three modulations (QAM4, QAM16 and QAM64) and two system sizes (32 and 16 transmitters, 64 receivers) with 3GPP MIMO channels. 36
- 4-3 SNR requirement gap with Maximum-Likelihood at SER of 10^{-3} . The total bar height shows the 90th-percentile gap (over different channels) while the hatched section depicts the average. 37

- 5-1 Noise power after the linear and denoiser stages at different layers of OAMP-Net and MMNet. The OAMPNet denoisers become ineffective after the third layer on 3GPP MIMO channels. 41
- 5-2 Percentage of transmitters that have Gaussian error distribution after the linear block for each layer with significance level of 5%. MMNet produces Normal-distributed errors at the output of linear blocks, while OAMPNet fails to achieve the Gaussian property. 42

5-3	Median of Anderson score for the noise at the output of linear stage vs. the power of noise at the input of the stage. MMNet controls the linear block output error distribution to be Gaussian. Dashed horizontal lines show the thresholds for 1%, 5% and 15 significance level.	43
5-4	Effect of condition number on performance of schemes. (a) MMNet is more robust in producing the right noise distribution for denoisers with changes in condition number. (b) SER is directly affected by the condition number.	45
6-1	Correlation of channel samples across time and frequency dimensions. The correlation decays relatively quickly in the time dimension, but the channel matrices show strong locality across sub-carriers in frequency dimension. .	49
6-2	CDF of SER using Algorithm 1 for training MMNet on QAM16 modulation. MMNet requires only 4 overall iterations of batch size 500 per channel realization to train to a reasonable performance.	51
6-3	Number of multiplication operations per signal detection for different algorithms on QAM16 with $N_r = 64$ receive antennas in 3GPP MIMO model. Detection with MMNet, including its online training process, requires fewer multiplication operations than detection with pre-trained DetNet and OAMPNet models.	51

Adaptive Neural Signal Detection for Massive MIMO

by

Mehrdad Khani Shirkoohi

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2019, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

Massive Multiple-Input Multiple-Output (MIMO) is a key enabler for fifth generation (5G) cellular communication systems. Massive MIMO gives rise to challenging signal detection problems for which traditional detectors are either impractical or suffer from performance limitations. Recent work has proposed several learning approaches to MIMO detection with promising results on simple channel models (e.g., i.i.d. Gaussian entries). However, we find that the performance of these schemes degrades significantly in real-world scenarios in which the channels of different receivers are spatially correlated. The root of this poor performance is that these schemes either do not exploit the problem structure (requiring models with millions of training parameters), or are overly-constrained to mimic algorithms that require very specific assumptions about the channel matrix.

We propose MMNet, a deep learning MIMO detection scheme that significantly outperforms existing approaches on realistic channel matrices with the same or lower computational complexity. MMNet's design builds on the theory of iterative soft-thresholding algorithms to identify the right degree of model complexity, and it uses a novel training algorithm that leverages temporal and frequency locality of channel matrices at a receiver to accelerate training. Together, these innovations allow MMNet to train online for every realization of the channel. On i.i.d. Gaussian channels, MMNet requires 2 orders of magnitude fewer operations than existing deep learning schemes but achieves near-optimal performance. On spatially-correlated realistic channels, MMNet achieves the same error rate as the next-best learning scheme (OAMPNet [1]) at 2.5dB lower Signal-to-Noise Ratio (SNR) and with at least $10\times$ less computational complexity. MMNet is also 4-8dB better overall than a classic linear scheme like the minimum mean square error (MMSE) detector.

Thesis Supervisor: Mohammad Alizadeh

Title: Assistant Professor of Electrical Engineering and Computer Science

Chapter 1

Introduction

The fifth generation of cellular communication systems (5G) promises an order of magnitude higher spectral efficiency (measured in bits/s/Hz) than legacy standards such as Long Term Evolution (LTE) [2]. One of the key enablers of this better efficiency is Massive Multiple-Input Multiple-Output (MIMO) [3], in which a base station (BS) equipped with a very large number of antennas (around 64–256) simultaneously serves multiple single-antenna user equipments (UEs) on the same time-frequency resource.

Legacy systems already use MIMO [4], but this is the first time it will be deployed on such a large scale, creating significant challenges for *signal detection*. The goal of signal detection is to infer the transmitted signal vector \mathbf{x} from the vector $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}$ received at the BS antennas, where \mathbf{H} is the channel matrix, and \mathbf{n} is Gaussian noise. Traditional MIMO detection methods with strong performance [5, 6, 7, 8] are feasible only for small systems and have prohibitive complexity for massive MIMO deployments. Thus, there is a need for low-complexity symbol detection schemes that perform well and scale to large system dimensions.

In recent work, researchers have proposed several learning approaches to MIMO detection. Samuel et al. [9] developed a deep neural network architecture called DetNet with impressive performance, e.g., matching the performance of a semidefinite relaxation (SDR)

baseline for i.i.d. Gaussian channel matrices while running $30\times$ faster. Shortly afterwards, inspired by the Orthogonal AMP algorithm [10], He et al. [1] introduced OAMPNet and demonstrated strong performance on both i.i.d. Gaussian and small-sized correlated channel matrices based on the Kronecker model with exponentially-distributed spatial correlations [11]. DetNet and OAMPNet are both trained offline: they try to learn a single model during training for a family of channel matrices (e.g., i.i.d. Gaussian channels). However, the two schemes have different design philosophies. DetNet embeds little domain knowledge into the model and relies on a large neural network with 1-10 million parameters depending on the system size and modulation scheme. By contrast, OAMPNet takes a model-driven approach and follows the OAMP algorithm closely; it adds only 2 trainable parameters per iteration of the OAMP algorithm.

In this dissertation we show that neither approach is effective in practice. We conduct extensive experiments using a dataset of channel realizations from the 3GPP 3D MIMO channel [12], as implemented in the QuaDRiGa channel simulator [13]. Our results show that DetNet’s training is unstable for realistic channels, while OAMPNet suffers a large performance gap (5–7dB at symbol error rate of 10^{-3}) compared to the optimal Maximum-Likelihood detector on these channels. Both models (as well as several classical baselines) perform well in simpler settings used for evaluation in prior work (e.g., i.i.d. Gaussian channels, low-order modulation schemes). Our results demonstrate the difficulty of learning a fixed detector that generalizes across a wide variety of channel matrices (esp. poorly-conditioned channels that are difficult to invert). DetNet’s approach is, in a sense, *too* general, making the large model difficult to train, while OAMPNet makes strong assumptions about channel matrices (OAMP was designed for unitarily-invariant channels [10]) and therefore performs poorly on channels that deviate from these assumptions.

Motivated by these findings, we revisit MIMO detection from an online learning perspective. We ask: *Can a receiver adapt its detector for every realization of the channel matrix?* Such an approach would arguably be simpler and could perform better than a fixed detector that must handle a wide variety of channel matrices. However, conventional wisdom suggests that training a deep neural network online is “impossible” in this context because

of the stringent performance requirements of MIMO detectors [9].

MMNet overcomes this challenge with two key ideas. First, it uses a neural network architecture that strikes the right balance between expressivity and complexity. MMNet’s neural network is based on iterative soft-thresholding algorithms [14, 15]. It preserves important aspects of these algorithms in MIMO detection, such as a denoiser architecture tailored for uncorrelated Gaussian noise for different transmitted signals. At the same, MMNet introduces adequate flexibility into these algorithms, with trainable parameters that are optimized for each channel realization. Second, MMNet’s online training algorithm exploits the locality of channel matrices at a receiver in both the frequency and time domains. By leveraging locality, MMNet accelerates training $250\times$ compared to naively retraining the neural network from scratch for each channel realization. Taken together, these ideas enable MMNet to achieve performance within 1.5dB of the optimal Maximum-Likelihood detector with $10\text{-}15\times$ less computational complexity than the second best scheme, OAMPNet. On random i.i.d. Gaussian channels, we show that a simple version of MMNet with $100\times$ less complexity than OAMPNet and DetNet, achieves near-optimal performance without requiring any retraining.

We empirically analyze the dynamics of errors across different layers of MMNet and OAMPNet to understand why MMNet achieves higher detection accuracy. Our analysis reveals that MMNet shapes the distribution of noise at the input of denoisers to ensure they operate effectively. In particular, as signals propagate through the MMNet neural network, the noise distribution at the input of denoiser stages approaches a Gaussian distribution, to create precisely the conditions in which the denoisers can attenuate noise maximally.

The rest of this dissertation is organized as follows. Chapter 2 provides background on classical and learning-based detection schemes, and introduces a general iterative framework that can express many of these algorithms. Chapter 3 introduces the MMNet design in addition to a simple variant for i.i.d. channels. Chapter 4 shows performance results of detection algorithms on i.i.d. Gaussian and 3GPP MIMO channels for different modulations. Chapter 5 discusses the error dynamics of MMNet and empirically studies why

it performs better than OAMPNet. Chapter 6 introduces the MMNet online training algorithm and how temporal and spectral locality of channel matrices can significantly reduce the computational complexity of training MMNet.

Notation: We will use lowercase symbols for scalars, bold lowercase symbols for vectors and bold uppercase symbols to denote matrices. Symbols $\{\theta, \boldsymbol{\theta}, \boldsymbol{\Theta}\}$ are used to represent the parameters of trainable models. The pseudo-inverse of the matrix \mathbf{H} is denoted by $\mathbf{H}^+ = (\mathbf{H}^H \mathbf{H})^{-1} \mathbf{H}^H$. \mathbf{I}_n stands for identity matrix of size n .

Chapter 2

Background and Related Work

This chapter introduces the MIMO detection problem and reviews the most relevant related work.

2.1 Problem Definition

We consider a communication channel from N_t single-antenna transmitters to a receiver equipped with N_r antennas. The received vector $\mathbf{y} \in \mathbb{C}^{N_r}$ is given as

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \quad (2.1)$$

where $\mathbf{H} \in \mathbb{C}^{N_r \times N_t}$ is the channel matrix, $\mathbf{n} \sim \mathcal{CN}(0, \sigma^2 \mathbf{I}_{N_r})$ is complex Gaussian noise and $\mathbf{x} \in \mathcal{X}^{N_t}$ is the vector of transmitted symbols. \mathcal{X} denotes the finite set of constellation points. We assume that each transmitter chooses a symbol from \mathcal{X} uniformly at random, and all transmitters use the same constellation set. Further, as is standard practice, we assume that the constellation set \mathcal{X} is given by a quadrature amplitude modulation (QAM) scheme [16]. All constellations are normalized to unit average power (e.g., the QAM4 constellation is $\{\pm \frac{1}{\sqrt{2}} \pm j \frac{1}{\sqrt{2}}\}$).

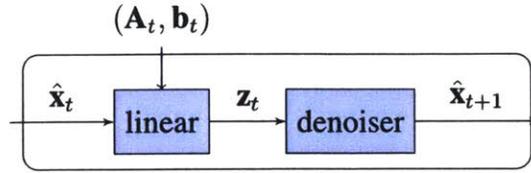


Figure 2-1: A block of an iterative detector in general framework. Each block contains a linear transformation followed by a denoising stage.

The channel matrix \mathbf{H} is generated by a stochastic process, but it is assumed to be perfectly known at the receiver. The goal of the receiver is to compute the maximum likelihood (ML) estimate $\hat{\mathbf{x}}$ of \mathbf{x} :

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathcal{X}^{N_t}} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2. \quad (2.2)$$

The optimization problem in Eq. (2.2) is NP-hard due to the finite-alphabet constraint $\mathbf{x} \in \mathcal{X}^{N_t}$ [17]. Over the last three decades, researchers have proposed a variety of detectors for this problem with differing levels of complexity. We briefly describe a small subset of existing detection schemes in this chapter. We refer the interested reader to [6, 7] for a comprehensive overview of MIMO detection schemes.

2.2 An iterative framework for MIMO detection

In this dissertation, we focus on a class of iterative estimation algorithms for solving Eq. (2.2) shown in Fig. 2-1. Each iteration of these algorithms comprises the following two steps:

$$\begin{aligned} \text{General Iteration:} \quad & \mathbf{z}_t = \hat{\mathbf{x}}_t + \mathbf{A}_t(\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}_t) + \mathbf{b}_t, \\ & \hat{\mathbf{x}}_{t+1} = \eta_t(\mathbf{z}_t). \end{aligned} \quad (2.3)$$

The first step takes as input $\hat{\mathbf{x}}_t$, a current estimate of \mathbf{x} and the received signal \mathbf{y} and applies a *linear* transformation to obtain an intermediate signal \mathbf{z}_t . In the second step, a *non-linear* “denoiser” is applied to \mathbf{z}_t to produce $\hat{\mathbf{x}}_{t+1}$, a new estimate of \mathbf{x} that is used as the input for the next iteration. Together, the linear and denoising operations aim to improve the quality

of the estimate $\hat{\mathbf{x}}_t$ from one iteration to the next.

We refer to $\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}_t$ as the *residual* term. The denoiser $\eta_t(\cdot)$ can be any non-linear function in general; however, most algorithms apply the same *thresholding* function $\beta_t: \mathbb{C} \rightarrow \mathbb{C}$ to each element. Using an element-wise thresholding function can significantly reduce the complexity of the denoising step. Typically denoisers also require one or more scalar parameters which depend on the detector information of the system (channel measurement, residuals, etc.) and they update by each iteration of the algorithm and we denote them by σ_t . We use the terms *step*, *layer* and *block* interchangeably to refer to one complete iteration (the linear step followed by the non-linear denoiser) of the algorithms in this dissertation. All algorithms discussed assume $\hat{\mathbf{x}}_0 = 0$.

A natural choice for the denoising function is the minimizer of $\mathbb{E}[\|\hat{\mathbf{x}} - \mathbf{x}\|_2 | \mathbf{z}_t]$, which is given by:

$$\eta(\mathbf{z}_t) = \mathbb{E}[\mathbf{x} | \mathbf{z}_t]. \quad (2.4)$$

Optimal denoiser for Gaussian noise

Assume that the noise at the input of the denoiser $\mathbf{z}_t - \mathbf{x}$ has an i.i.d. Gaussian distribution with diagonal covariance matrix $\sigma_t^2 \mathbf{I}_{N_t}$. The element-wise thresholding function derived from Eq. (2.4) has the form

$$\beta_t^g(z; \sigma_t^2) = \frac{1}{Z} \sum_{x_i \in \mathcal{X}} x_i \exp\left(-\frac{\|z - x_i\|^2}{\sigma_t^2}\right), \quad (2.5)$$

where $Z = \sum_{x_j \in \mathcal{X}} \exp\left(-\frac{\|z - x_j\|^2}{\sigma_t^2}\right)$. As we see here, σ_t in Eq. (2.5) denoiser represents the standard deviation of the Gaussian noise on denoiser inputs. In all denoisers in $\eta_t(\cdot; \sigma_t^2)$ format in this dissertation, σ_t^2 refers to the variance of noise in denoiser input.

In the following, we briefly describe several algorithms for MIMO detection. We begin with traditional, non-learning approaches (§ 2.3) and then discuss recent deep learning proposals (§ 2.4). We show how many of these algorithms can be expressed in the iterative

framework discussed above.

2.3 Classical MIMO detection algorithms

2.3.1 Linear

The simplest method to approximately solve Eq. (2.2) is to relax the constraint of $\mathbf{x} \in \mathcal{X}^{N_t}$ to $\mathbf{x} \in \mathbb{C}^{N_t}$ and then round the relaxed solution to a point on the constellation:

$$\begin{aligned} \text{Linear:} \quad \mathbf{z} &= \arg \min_{\mathbf{x} \in \mathbb{C}^{N_t}} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2 = \mathbf{H}^+ \mathbf{y}, \\ \hat{\mathbf{x}} &= \arg \min_{\mathbf{x} \in \mathcal{X}^{N_t}} \|\mathbf{x} - \mathbf{z}\|_2. \end{aligned} \tag{2.6}$$

Rounding each component of \mathbf{z} to the closest point in the constellation set $\hat{\mathbf{x}}$ leads to the well-known zero-forcing (ZF) detector, which is equivalent to a single step of Eq. (2.3) with initial condition of $\hat{\mathbf{x}}_0 = 0$, $\mathbf{A}_0 = \mathbf{H}^+$, $\mathbf{b}_0 = 0$ and a hard-decision denoiser with respect to the points in the constellation. Other widely-used single-step linear detectors include the matched filter and the minimum mean square error (MMSE) detectors [3] with $\mathbf{A}_0 = \mathbf{H}^H$ and $\mathbf{A}_0 = (\mathbf{H}^H \mathbf{H} + \sigma^2 \mathbf{I}_{N_t})^{-1} \mathbf{H}^H$, respectively. Linear detectors are attractive for practical implementation because of their low complexity, but they perform substantially worse than the optimal detector.

We can also perform the optimization in Eq. (2.6) in multiple iterations using gradient descent. The gradient of the objective function in the first equation of Eq. (2.6) with respect to \mathbf{x} is $-2\mathbf{H}^H(\mathbf{y} - \mathbf{H}\mathbf{x})$. Hence, if we set \mathbf{A}_t to $2\alpha\mathbf{H}^H$ and $\mathbf{b}_t = 0$, the linear step of Eq. (2.3) is equivalent to minimizing $\|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2$ using gradient descent with step size α . This is followed by a mapping onto the constellation set in the denoising step. If we had a compact convex constellation set, this projected gradient descent procedure is guaranteed to converge to the global optimum. Discrete constellation sets, however, are not compact convex. Nonetheless, solving the linear least squares problem in Eq. (2.6) iteratively may be desirable to avoid the cost of calculation the pseudo-inverse of the channel matrix.

2.3.2 Approximate Message Passing (AMP)

MIMO detection can, in principle, be solved through belief propagation (BP) if we consider a bipartite graph representation of the model in Eq. (2.1) [18]. BP on this graph requires $\mathcal{O}(N_r N_t)$ update messages in each iteration, which would be prohibitive for large system dimensions. In the large system limit, Jeon et al. [15] introduce approximate message passing (AMP) as a lower complexity inference algorithm for solving Eq. (2.2) for i.i.d. Gaussian channels. AMP reduces the number of messages in each iteration to $\mathcal{O}(N_r + N_t)$. The algorithm performs the following sequence of updates:

$$\begin{aligned} \mathbf{z}_t &= \hat{\mathbf{x}}_t + \mathbf{H}^H(\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}_t) + \mathbf{b}_t, \\ \text{AMP:} \quad \mathbf{b}_t &= \alpha_t (\mathbf{H}^H(\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}_{t-1}) + \mathbf{b}_{t-1}), \\ \hat{\mathbf{x}}_{t+1} &= \eta_t(\mathbf{z}_t; \sigma_t). \end{aligned} \tag{2.7}$$

Consider AMP in our iterative framework: we use $\mathbf{A}_t = \mathbf{H}^H$ as the linear operator. The vector \mathbf{b}_t is known as the *Onsager* term. The scalar sequences σ_t and α_t can be computed given the SNR and system parameters (constellation and number of transmitters and receivers) [19]. The denoising function $\eta_t(\cdot)$ applies the optimal denoiser for Gaussian noise in Eq. (2.5) to each element of the vector \mathbf{z}_t . Jeon et al. [15] prove that AMP is asymptotically optimal for large i.i.d. Gaussian channel matrices.

Orthogonal AMP (OAMP) [10] was proposed for unitarily-invariant channel matrices [20] to relax the i.i.d. Gaussian channel assumption in the original AMP algorithm.

$$\begin{aligned} \text{OAMP:} \quad \mathbf{z}_t &= \hat{\mathbf{x}}_t + \gamma_t \mathbf{H}^H (v_t^2 \mathbf{H} \mathbf{H}^H + \sigma^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}_t), \\ \hat{\mathbf{x}}_{t+1} &= \eta_t(\mathbf{z}_t; \sigma_t^2). \end{aligned} \tag{2.8}$$

In Eq. (2.8), $\gamma_t = N_t / \text{trace} \left(v_t^2 \mathbf{H}^H (v_t^2 \mathbf{H} \mathbf{H}^H + \sigma^2 \mathbf{I})^{-1} \mathbf{H} \right)$ is a normalizing factor and v_t^2 is proportional to the average noise power at the denoiser output at iteration t and can be computed given the SNR and system dimensions. Notice that OAMP requires computing a matrix inverse in each iteration, making it more computationally expensive than

AMP.

2.3.3 Other techniques

Several detection schemes relax the lattice constraint ($\mathbf{x} \in \mathcal{X}^{N_t}$) in Eq. (2.2). For example, Semi-Definite Relaxation (SDR) [8] formulates the problem as a semi-definite program. Sphere decoding [5] conducts a search over solutions $\hat{\mathbf{x}}$ such that $\|\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}\|_2 \leq r$. Increasing r covers a larger set of possible solutions, but this comes at the cost of increased complexity, approaching that of brute-force search. There is a large body of variations on improvements to this idea, which can be found in [6, 7]. While these approaches can perform well, their computational complexity is prohibitive for Massive MIMO systems with currently available hardware.

Another class of detector applies several stages or iterations of linear detection followed by interference subtraction from the observation \mathbf{y} . The V-BLAST scheme [21] does this by detecting the strongest symbols, which are then successively removed from \mathbf{y} . The drawbacks of this approach are error propagation of early symbol decisions and high complexity due to the N_t required stages, as well as the necessary reordering of transmitters after each step. Parallel interference cancellation (PIC) has been proposed to circumvent these problems. PIC jointly detects all transmitted symbols and then attempts to create an interference-free channel for each transmitter through the cancellation of all other transmitted symbols [22, 23]. A large system approximation of this approach was recently developed in [24] based on [25]. However, it is currently limited to binary phase shift keying (BPSK) modulation and leads to unsatisfactory performance for realistic system dimensions.

In summary, most existing techniques are too complex to implement at the scale required by next-generation Massive MIMO systems. On the other hand, light-weight techniques like AMP cannot handle correlated channel matrices. These limitations have motivated a number of learning-based proposals for MIMO detection, which we discuss next.

2.4 Learning-based MIMO detection schemes

2.4.1 DetNet

Inspired by iterative projected gradient descent optimization, Samuel et al. [26] propose DetNet, a deep neural network architecture for MIMO detection. This architecture performs very well in cases of i.i.d. complex Gaussian channel matrices and achieves the performance of state-of-the-art algorithms in lower-order modulations, such as BPSK and QAM4. However, it is far more complex. The neural network is described by the following equations:

$$\begin{aligned}
 \mathbf{q}_t &= \hat{\mathbf{x}}_{t-1} - \theta_t^{(1)} \mathbf{H}^H \mathbf{y} + \theta_t^{(2)} \mathbf{H}^H \mathbf{H} \hat{\mathbf{x}}_{t-1}, \\
 \mathbf{u}_t &= \left[\Theta_t^{(3)} \mathbf{q}_t + \Theta_t^{(4)} \mathbf{v}_{t-1} + \theta_t^{(5)} \right]_+, \\
 \mathbf{v}_t &= \Theta_t^{(6)} \mathbf{u}_t + \theta_t^{(7)}, \\
 \hat{\mathbf{x}}_t &= \Theta_t^{(8)} \mathbf{u}_t + \theta_t^{(9)}.
 \end{aligned} \tag{2.9}$$

where $[x]_+ = \max(x, 0)$, which is also known as ReLU activation function [27], is applied element-wise.

Although DetNet's performance is promising, it has two main limitations. First, its heuristic nature makes it difficult to reason about how the neural network works, and how to extend its architecture, for example, to support spatially correlated channel matrices or higher-order modulation schemes. Second, DetNet's architecture does not incorporate known properties of such iterative methods, and thus it is unnecessarily complex. For example, many iterative soft-thresholding schemes (including AMP described above) apply a denoiser tailored for Gaussian noise (Eq. (2.5)) independently to each transmitted signal. DetNet's neural network can also be thought to be performing non-linear denoising steps intermixed with linear transformations. However, DetNet's denoisers are fully-connected 2-layer neural networks that operate on the entire vector of transmitted signals, i.e., they are N_t -dimensional functions instead of simple scalar functions.

2.4.2 OAMPNet

He et al. [1] design a learning-based iterative scheme based on the OAMP algorithm. OAMPNet adds two tuning parameters per iteration to the OAMP algorithm. OAMPNet shows very good performance in the case of i.i.d. Gaussian channels, but it does not generalize to realistic channels with spatial correlations, as our experiments in Chapter 4 show. The OAMPNet design can be expressed as:

$$\begin{aligned} \text{OAMPNet:} \quad \mathbf{z}_t &= \hat{\mathbf{x}}_t + \theta_t^{(1)} \mathbf{H}^H (v_t^2 \mathbf{H} \mathbf{H}^H + \sigma^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{H} \hat{\mathbf{x}}_t) \\ \hat{\mathbf{x}}_{t+1} &= \eta_t (\mathbf{z}_t; \sigma_t^2). \end{aligned} \quad (2.10)$$

OAMPNet uses the same denoisers used by AMP, which are optimal for Gaussian noise.

By basing its design on OAMP, OAMPNet makes a strict assumption about the system: unitarily-invariant channel matrices. This reduces OAMPNet to training a few parameters per iteration. However, as our results show, OAMPNet's assumptions make it brittle, and its performance degrades on real channel matrices that do not conform to the assumptions. Further, like OAMP, OAMPNet must compute a matrix pseudo-inverse in each iteration, and therefore its complexity is still quite high compared to schemes like AMP.

Chapter 3

MMNet Design

The MMNet design follows the iterative framework described in § 2.2. The main idea behind MMNet is to introduce the *right degree of flexibility* into the linear and denoising components of the iterative framework while preserving its overall structure. We observe that prior architectures do not strike the right balance between model flexibility and complexity. DetNet is very general and expressive, but it requires an enormous number of parameters because it does not incorporate domain knowledge, such as the role and structure of the denoiser in iterative soft-thresholding algorithms. On the other hand, OAMPNet is heavily constrained and follows OAMP very closely; however, OAMP makes rigid assumptions about channel matrices (unitarily-invariant channels). Thus, as our results will show, OAMPNet’s performance degrades on real channel matrices that violate this assumption.

In practice, not much is known about the distribution of channel matrices. We therefore propose a data-driven approach in which we learn a set of model parameters for each realization of \mathbf{H} . In this approach, the receiver continually adapts its parameters as it measures new channel matrices \mathbf{H} . We demonstrate that this can be realized in practice with a suitable neural network architecture by exploiting the fact that realistic channels exhibit locality in both the frequency and time domains. We introduce the neural network architecture in this chapter and discuss a practical training algorithm in Chapter 6.

We present separate neural network models for (1) i.i.d. Gaussian channels and (2) arbitrary channels. In the i.i.d. Gaussian case, the model is extremely simple:

$$\begin{aligned} \text{MMNet-iid:} \quad \mathbf{z}_t &= \hat{\mathbf{x}}_t + \theta_t^{(1)} \mathbf{H}^H (\mathbf{y} - \mathbf{H} \hat{\mathbf{x}}_t), \\ \hat{\mathbf{x}}_{t+1} &= \eta_t (\mathbf{z}_t; \sigma_t^2). \end{aligned} \quad (3.1)$$

Here, the denoiser is the optimal denoiser for Gaussian noise given in Eq. (2.5). MMNet-iid assumes the same distribution of noise at the input of the denoiser for all transmitted symbols and estimates its variance σ_t^2 according to

$$\sigma_t^2 = \frac{\theta_t^{(2)}}{N_t} \left(\frac{\|\mathbf{I} - \mathbf{A}_t \mathbf{H}\|_F^2}{\|\mathbf{H}\|_F^2} [\|\mathbf{y} - \mathbf{H} \hat{\mathbf{x}}_t\|_2^2 - N_r \sigma^2]_+ + \frac{\|\mathbf{A}_t\|_F^2}{\|\mathbf{H}\|_F^2} \sigma^2 \right). \quad (3.2)$$

The intuition behind Eq. (3.2) is that the noise at the input of the denoiser at step t is comprised of two parts: (1) the residual error caused by deviation of $\hat{\mathbf{x}}_t$ from the true value of \mathbf{x} , and (2) the contribution of the channel noise \mathbf{n} . The first component is amplified by the linear transformation $(\mathbf{I} - \mathbf{A}_t \mathbf{H})$, and the second component is amplified by \mathbf{A}_t . See [10, 15] for further details on this method for estimating noise variance.

This model has only two parameters per layer: $\theta_t^{(1)}$ and $\theta_t^{(2)}$. We discuss this model merely to illustrate that, for the i.i.d. Gaussian channel matrix case (which most prior work has used for evaluation), a simple model that adds a small amount of flexibility to existing algorithms like AMP can perform very well. In fact, our results will show that, in this case, we do not even need to train the parameters of the model online for each channel realization; training offline over randomly sampled i.i.d. Gaussian channel suffices.

The MMNet neural network for arbitrary channel matrices is as follows:

$$\begin{aligned} \text{MMNet:} \quad \mathbf{z}_t &= \hat{\mathbf{x}}_t + \Theta_t^{(1)} (\mathbf{y} - \mathbf{H} \hat{\mathbf{x}}_t), \\ \hat{\mathbf{x}}_{t+1} &= \eta_t (\mathbf{z}_t; \sigma_t^2). \end{aligned} \quad (3.3)$$

In Eq. (3.3), $\Theta_t^{(1)}$ is an $N_t \times N_r$ complex-valued trainable matrix. In order to enable the model to handle cases in which different transmitted symbols have differing levels of noise,

we add an extra degree of freedom to our estimations of noise per transmitter, resulting in:

$$\sigma_t^2 = \frac{\boldsymbol{\theta}_t^{(2)}}{N_t} \left(\frac{\|\mathbf{I} - \mathbf{A}_t \mathbf{H}\|_F^2}{\|\mathbf{H}\|_F^2} [\|\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}_t\|_2^2 - N_r \sigma^2]_+ + \frac{\|\mathbf{A}_t\|_F^2}{\|\mathbf{H}\|_F^2} \sigma^2 \right). \quad (3.4)$$

Parameter vector $\boldsymbol{\theta}_t^{(2)}$ of size $N_t \times 1$ scales the noise variance by different amounts for each symbol. This approach distinguishes MMNet from both the highly-constrained OAMPNet and overly complex DetNet solution. In particular, MMNet uses a flexible linear transformation (which does not need to be linear in \mathbf{H}) to construct the intermediate signal z_t , but it applies the standard optimal denoiser for Gaussian noise in Eq. (2.5). Further, unlike OAMPNet, MMNet does not require any expensive matrix inverse operation.

MMNet concatenates T layers of the above form. We use the the average L2-loss over all T layers in order to train the model. This loss is denoted by L in Eq. (3.5).

$$L = \frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{x}}_t - \mathbf{x}\|_2^2. \quad (3.5)$$

Chapter 4

Experiments

In this chapter, we evaluate and compare the performance of MMNet with state-of-the-art schemes for both i.i.d. Gaussian and realistic channel matrices. These are our main findings:

1. On i.i.d. Gaussian channels, most schemes perform very well. MMSE, SDR, V-BLAST and DetNet are 1-2dB far from the best schemes overall. AMP performance degrades for higher-order modulations in high SNRs. MMNet-iid and OAMPNet are very close to Maximum-Likelihood in all experiment on these channels. MMNet-iid, however, has 2 orders of magnitude lower complexity than learning-based schemes OAMPNet and DetNet.
2. On realistic, spatially-correlated channel matrices, the performance of all existing learning-based approaches degrades significantly. MMNet shows the least gap with Maximum-Likelihood in all schemes. While DetNet and AMP fail to extend to these channels with reasonable performance (on QAM4, for example), MMSE has an 8-10dB gap with Maximum-Likelihood on 64×16 channels. OAMPNet reduces this gap to 5-7dB. However, MMNet closes the gap to less than 1.5dB.

4.1 Methodology

We first briefly discuss the details of detection schemes used for comparison. Since some of these schemes (including MMNet) require training, we then discuss the process of generating data and training/testing on this data.

4.1.1 Compared Schemes

In our experiments, we compare the following schemes on QAM modulation:

- **MMSE:** Linear decoder that applies the SNR-regularized channel's pseudo inverse and rounds the output to the closest point on the constellation (see § 2.3.1).
- **SDR:** Semidefinite programming using a rank-1 relaxation interior point method [28].
- **V-BLAST:** Multi-stage interference cancellation BLAST algorithm using Zero-Forcing as the detection stage introduced in [23].
- **AMP:** AMP algorithm for MIMO detection from Jeon et al. [15]. AMP runs 50 iterations of the updates as discussed in Eq. (2.7). We verified that adding more iterations does not improve the results.
- **DetNet:** The deep learning approach introduced in [26]. The DetNet paper describes instantiations of the architecture for BPSK, QAM4 and QAM16; these neural networks have, on the order of 1-10M, trainable parameters depending on the size of the system and constellation set.
- **OAMPNet:** The OAMP-based architecture [1] implemented in 10 layers with 2 trainable parameters per layer and an inverse matrix computation at each layer.
- **MMNet-iid:** The simple architecture described in Chapter 3. This scheme has only 2 scalar parameters per layer and does not require any matrix inversions. We implement this neural network with 10 layers.

- **MMNet:** Our design described in Eq. (3.3). It has 10 blocks, and the total number of trainable parameters is $2N_t(N_r + 1)$ real values, independent of constellation size.
- **Maximum-Likelihood:** The optimal solver for Eq. (2.2) using a highly-optimized Mixed Integer Programming package Gurobi [29].

4.1.2 Dataset

Training and test data are generated through the model described in Eq. (2.1). In this model, there are three sources of randomness: signal \mathbf{x} , channel noise \mathbf{n} and channel matrix \mathbf{H} . Each transmitted signal \mathbf{x} is generated randomly and uniformly over the corresponding constellation set. All transmitters are assumed to use the same modulation. The channel noise \mathbf{n} is sampled from a zero-mean i.i.d. normal distribution with a variance that is set according to the operating SNR, defined as $\text{SNR}(\text{dB}) = 10 \log (\mathbb{E}[\|\mathbf{H}\mathbf{x}\|_2^2] / \mathbb{E}[\|\mathbf{n}\|_2^2])$. For every training batch, the SNR(dB) is chosen uniformly at random in the desired operating SNR interval. This interval depends on the modulation scheme. For each modulation in each experiment, the SNR regime is chosen such that the best scheme other than Maximum-Likelihood can achieve a SER of 10^{-3} – 10^{-2} .

The channel matrices \mathbf{H} are either sampled from an i.i.d. Gaussian distribution (i.e., each column of \mathbf{H} is a complex-normal $\mathcal{CN}(0, (1/N_r)\mathbf{I}_{N_r})$), or they are generated via the realistic channel simulation described below.

Here, we study two system size ratios (N_t/N_r) of 0.25 and 0.5, with the total number of receivers fixed at $N_r = 64$. These are typical values for 4G/5G base stations in urban cellular deployments. For the case of realistic channels, we generate a dataset of channel realizations from the 3GPP 3D MIMO channel model [12], as implemented in the QuaDRiGa channel simulator [13].¹ We consider a base station (BS) equipped with a rectangular planar array consisting of 32 dual-polarized antennas installed at a height of 25 m. The BS is assumed to cover a 120°-cell sector of radius 500 m within which $N_t \in \{16, 32\}$ single-

¹The simulation results were generated using QuaDRiGa Version 2.0.0-664.

antenna users are dropped randomly. A guard distance of 10 m from the BS is kept. Each user is then assumed to move along a linear trajectory with a speed of 1 m/s. Channels are sampled every $\lambda/4$ m at a center frequency of 2.53 GHz to obtain sequences of length 100. Each channel realization is then converted to the frequency domain assuming a bandwidth of 20 MHz and using 1024 sub-carriers from which only every fourth is kept, resulting in $F = 256$ effective sub-carriers. We gather a total of 40 user drops, resulting in 40×256 length 100 sequences of channel matrices (i.e., 1M channel matrices in total). Since the pathloss can vary dramatically between different users, we assume perfect power control, which normalizes the average received power across antennas and sub-carriers to one. Denote $\mathbf{H}[f, k]$ as the k th column of channel matrix on sub-carrier f . Our normalization ensures that

$$\frac{1}{N_r N_t F} \sum_{k=1}^{N_t} \sum_{f=1}^F \|\mathbf{H}[f, k]\|^2 = 1.$$

4.1.3 Training

MMNet, DetNet, and OAMPNet require training and were implemented in TensorFlow [30]. We have converted Eq. (2.2) to its equivalent real-valued representation for TensorFlow implementations (cf. [31, Sec. II]). DetNet and OAMPNet are both trained as described in the corresponding publications (i.e., batch size of 5K samples). We trained each of the latter two algorithms for 50K iterations.

To train MMNet, we use the Adam optimizer [32] with a learning rate of 10^{-3} . Each training batch has a size of 500 samples. We train MMNet for 1K iterations on each realization of \mathbf{H} in the naive implementation. In § 6.2, we exploit frequency and time domain correlations to reduce the training requirement to 4 iterations per channel matrix.

For i.i.d. Gaussian channels, MMNet-iid is not trained per channel realization \mathbf{H} . Instead, we use 10K iterations with a batch size of 500 samples to train a single MMNet-iid neural network, which we then test on new channel samples.

4.2 Results

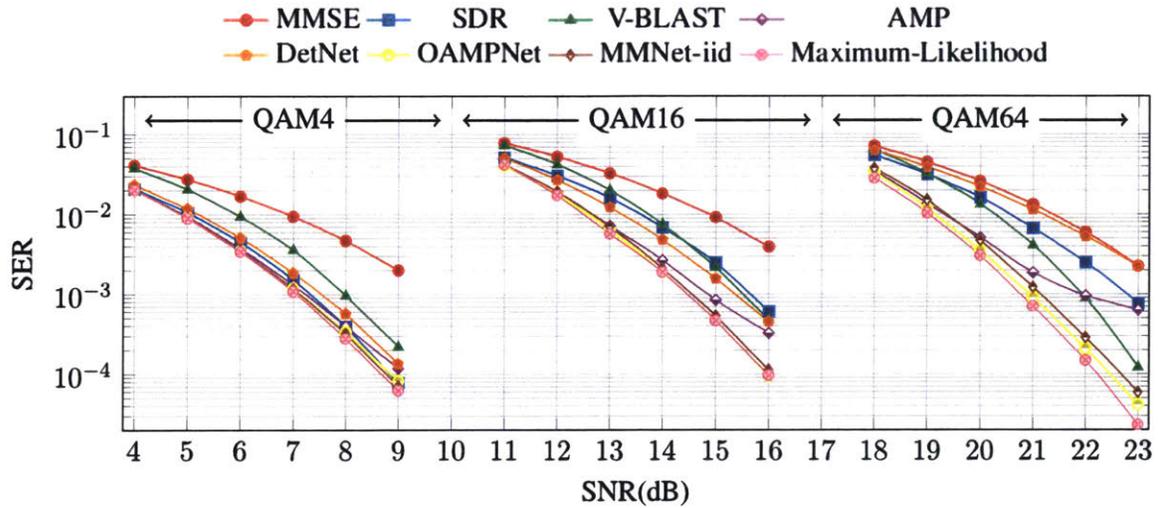
We compare different schemes along two axes: performance and complexity. In this section, we focus on performance, leaving a comparison of complexity to Chapter 6. We use the SNR required to achieve an SER of 10^{-3} as the primary performance metric. In practice, most error correcting schemes operate around a SER of 10^{-3} – 10^{-2} , so this is the relevant regime for MIMO detection.

4.2.1 i.i.d. Gaussian channels

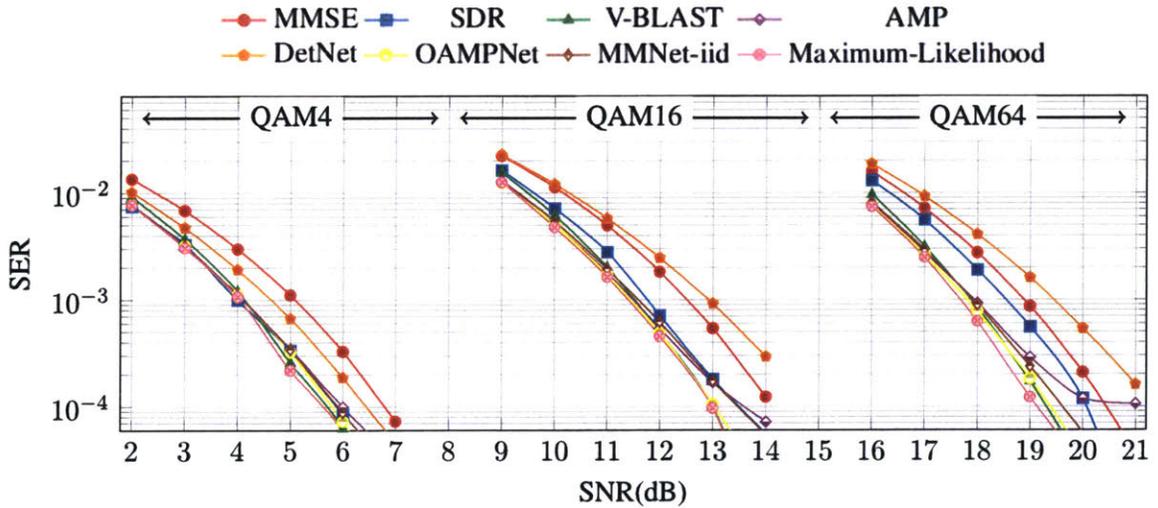
Fig. 4-1 shows the SER vs. SNR of the state-of-the-art MIMO schemes on Complex-Gaussian i.i.d. channels for two system sizes: 32 and 16 transmitters (Fig. 4-1a and Fig. 4-1b, respectively).

We make the following observations:

1. The SNR required to achieve a certain SER increases by 2–3dB as we double the number of transmitters (notice the different range of x-axes). The reason is that there is higher interference with more transmitters.
2. There is a 2–3dB performance gap between Maximum-Likelihood and MMSE across all modulations for $N_t = 32$. However, this gap decreases to 1dB for $N_t = 16$, because received signals have lower interference in this case.
3. Multiple schemes perform similarly to Maximum-Likelihood, especially at lower-order modulations like QAM4. As we move to QAM64, the performance of several schemes degrades compared to Maximum-Likelihood.
4. SDR performs better than MMSE, but its gap with Maximum-Likelihood increases with modulation order.
5. V-BLAST achieves almost the optimal performance across all modulations when we



(a) 32 Transmit, 64 Receive Antennas



(b) 16 Transmit, 64 Receive Antennas

Figure 4-1: SER vs. SNR of different schemes for three modulations (QAM4, QAM16 and QAM64) and two system sizes (32 and 16 transmitters, 64 receivers) with i.i.d. Gaussian channels.

have 16 transmit antennas. However, its performance is sensitive to system size and degrades when we increase the number of transmitters to 32.

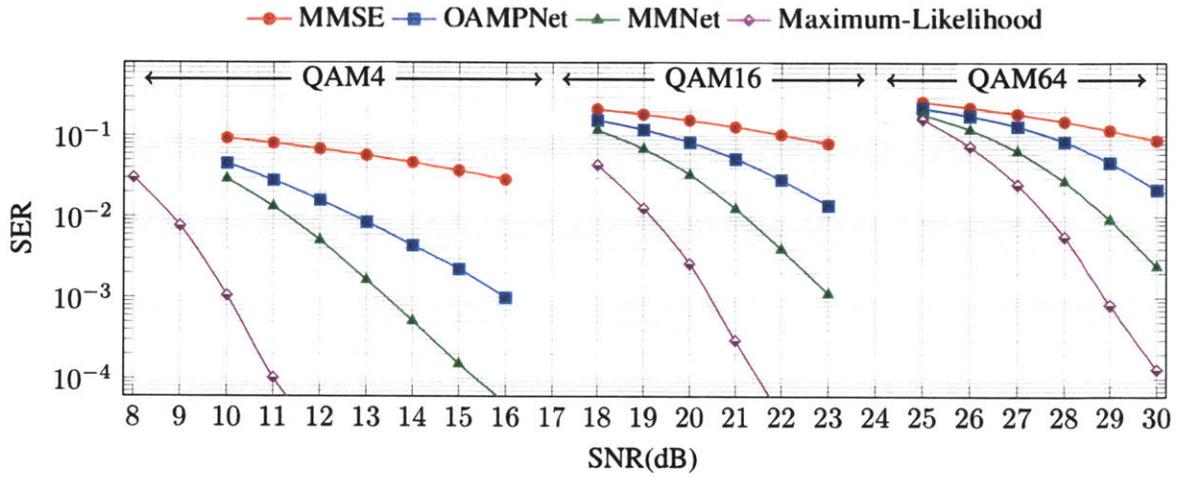
6. AMP is near-optimal in many cases (recall that, theoretically, AMP is asymptotically optimal for i.i.d. Gaussian channels). However, it suffers from robustness issues at higher SNR levels, especially with higher-order modulations like QAM64.
7. DetNet has a good performance on QAM4, but its gap with Maximum-Likelihood increases as we move to QAM16 and QAM64. With QAM64, DetNet performs even worse than MMSE for $N_t = 16$.
8. OAMPNet and the simple MMNet-iid approach are both very close to Maximum-Likelihood across different modulations over a wide range of SNRs, even though these models have only two parameters per layer. Unlike OAMPNet, MMNet-iid does not require any matrix inversions. As we discuss in Chapter 6, MMNet-iid has $\mathcal{O}(N_r) \times$ lower computational complexity than OAMPNet because it does not need matrix inversions and must learn only 20 parameters in total, compared to the more than 1M trainable parameters of DetNet.

In summary, these results show that for i.i.d. Gaussian channel matrices, adding just a small amount of flexibility via tuning parameters to existing iterative schemes like AMP can result in equivalent or improved performance over much more complex deep learning models like DetNet. These models may even outperform classical algorithms like SDR.

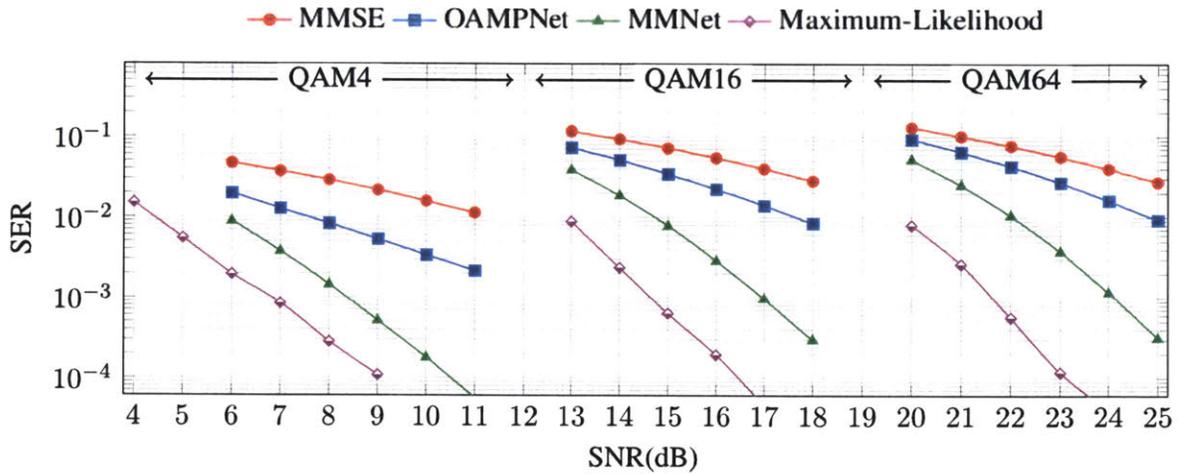
4.2.2 Realistic channels

Fig. 4-2 shows the results for the realistic channels derived using the 3GPP 3D MIMO channel model. We consider only MMSE (as a baseline), OAMPNet, MMNet and Maximum-Likelihood. As shown in the i.i.d. Gaussian case, schemes like SDR, V-BLAST and DetNet do not perform as well as the OAMPNet baseline.² Also, AMP is not designed for corre-

²We tried to train DetNet for realistic channels and ran into significant difficulty with stability and convergence in training.



(a) 32 Transmit, 64 Receive Antennas



(b) 16 Transmit, 64 Receive Antennas

Figure 4-2: SER vs. SNR of different schemes for three modulations (QAM4, QAM16 and QAM64) and two system sizes (32 and 16 transmitters, 64 receivers) with 3GPP MIMO channels.

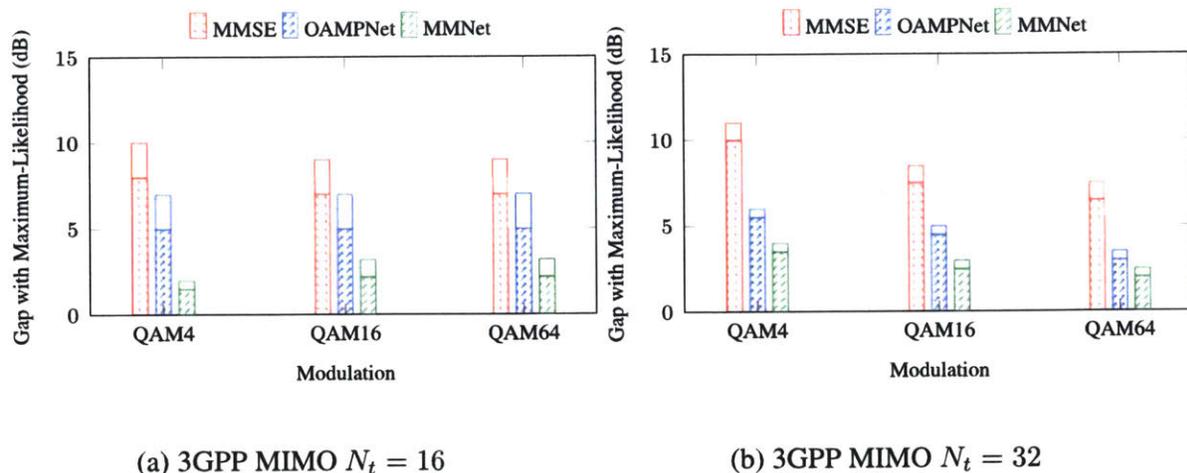


Figure 4-3: SNR requirement gap with Maximum-Likelihood at SER of 10^{-3} . The total bar height shows the 90th-percentile gap (over different channels) while the hatched section depicts the average.

lated channels and is known to perform poorly (see discussion in chapter 5).

We make the following observations:

1. There is a much larger gap with Maximum-Likelihood for all detection schemes on these channels compared to the i.i.d. case.
2. We require 4-7dB increase in the SNR ranges relative to Fig. 4-1. Also, doubling the number of transmitters from Fig. 4-2b to Fig. 4-2a costs about 5dB in SNR for each scheme this time (compare with 2-3dB in i.i.d. case.)
3. MMSE as a baseline shows a relatively flat SER vs. SNR in this case. For example, it requires 5dB SNR improvement on QAM16 to go from an SER of 2% to 1%.
4. OAMPNet performance improvement slope is faster than MMSE. It shows 2-3dB average improvement in SNR requirement relative to MMSE to achieve the same SER.
5. MMNet outperforms MMSE and OAMPNet schemes for both system sizes and in all modulations.

In Fig. 4-3, we plot the performance gap with Maximum-Likelihood for these three detec-

tion schemes. For this purpose, we measure the difference in the minimum SNR level that is required to have SER of 10^{-3} . In this figure, we have also attempted to show the variability of this requirement across different channel situations by adding the 90th-percentile gap over different channel conditions. We observe that MMNet can achieve up to 5dB and 8dB improvement, respectively, over OAMPNet and MMSE on more realistic channels.

Chapter 5

Why MMNet works

In this chapter, we examine why MMNet performs better than other schemes. By analyzing the dynamics of the error ($\hat{\mathbf{x}}_t - \mathbf{x}$), we find that MMNet’s denoisers are significantly more effective than those in OAMPNet. We show that this occurs because MMNet’s linear stages control the distribution of noise at the input of the denoisers. Specifically, MMNet ensures that the noise input to the denoisers is nearly Gaussian, whereas the noise distribution for OAMPNet is far from Gaussian. Since the denoisers in both architectures are tailored for Gaussian noise, they perform much better in MMNet.

5.1 Error dynamics

Define the error at the outputs of the linear and denoiser stages at iteration t as $\mathbf{e}_t^{lin} = \mathbf{z}_t - \mathbf{x}$ and $\mathbf{e}_t^{den} = \hat{\mathbf{x}}_{t+1} - \mathbf{x}$, respectively. For algorithms such as MMNet and OAMPNet where $\mathbf{b}_t = 0$, We can rewrite the update equations of Eq. (2.3) in terms of these two errors in the form:

$$\mathbf{e}_t^{lin} = (\mathbf{I} - \mathbf{A}_t \mathbf{H}) \mathbf{e}_{t-1}^{den} + \mathbf{A}_t \mathbf{n}, \quad (5.1a)$$

$$\mathbf{e}_t^{den} = \eta_t (\mathbf{x} + \mathbf{e}_t^{lin}) - \mathbf{x}. \quad (5.1b)$$

Eq. (5.1a) includes two terms, corresponding to the contribution of the error of the previous stages' output and the channel noise to \mathbf{e}_t^{lin} respectively. To gain intuition, consider the effect of several choices of \mathbf{A}_t .

If we set \mathbf{A}_t to \mathbf{H}^+ (the pseudo-inverse of the channel matrix), we are only left with the term $\mathbf{H}^+ \mathbf{n}$ in \mathbf{e}_t^{lin} , thus eliminating the error caused by the previous stage entirely. However, this comes at a price: we are left with Gaussian noise with covariance matrix $\sigma^2 \mathbf{H}^+ \mathbf{H}^{+H}$. This presents two potential problems: (1) if \mathbf{H} is ill-conditioned, it might amplify the channel noise (e.g., inversely proportional to the smallest singular value of \mathbf{H} in some directions); (2) if \mathbf{e}_t^{lin} is correlated noise, applying an element-wise denoising function to it may not be effective.

We could also remove the channel noise term entirely by setting \mathbf{A}_t to zero. This would effectively remove the linear stage. However, if optimal denoisers are used, removing the linear stage and applying the denoiser function twice should have no effect on reducing the error.

For i.i.d. Gaussian channels with variance $1/N_r$, if we set $\mathbf{A}_t = \mathbf{H}^H$, the factor $\mathbf{I} - \mathbf{A}_t \mathbf{H}$ asymptotically goes to zero as we increase N_r [19]. The auto-covariance of $\mathbf{A}_t \mathbf{n}$, $\sigma^2 \mathbf{H}^H \mathbf{H}$, is approximately equal to $\sigma^2 \mathbf{I}_{N_t}$. This means that the channel noise term is neither amplified nor correlated following this linear transformation with i.i.d. channels, while the error from the previous stage, \mathbf{e}_{t-1}^{den} , is attenuated significantly via $\mathbf{I} - \mathbf{A}_t \mathbf{H}$. These calculations explain why AMP has great performance on i.i.d. Gaussian channels. However, in case of correlated channels, neither $\mathbf{I} - \mathbf{A}_t \mathbf{H}$ is close to zero, nor is $\mathbf{A}_t \mathbf{n}$ uncorrelated. This is the primary reason that AMP cannot perform well on more realistic channels.

5.2 Analysis

As noted earlier, the key element in MMNet and prior schemes such as OAMPNet is how to pick the linear transformation \mathbf{A}_t . Based on the above discussion of the error dynamics,

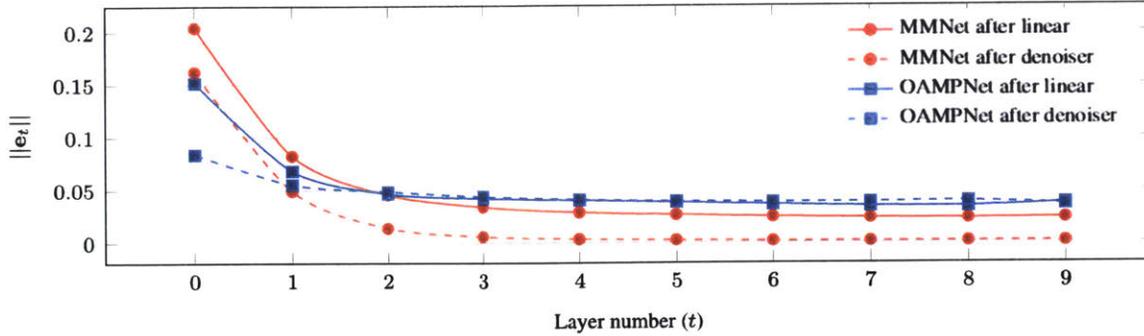


Figure 5-1: Noise power after the linear and denoiser stages at different layers of OAMPNet and MMNet. The OAMPNet denoisers become ineffective after the third layer on 3GPP MIMO channels.

we identify two desirable properties:

1. \mathbf{A}_t must reduce the magnitude of \mathbf{e}_t^{lin} . This requires striking a *balance* between the two terms in Eq. (5.1a), because attenuating one term may amplify the other and vice-versa.
2. \mathbf{A}_t must “shape” the distribution of \mathbf{e}_t^{lin} to make it suitable for the subsequent denoiser. In particular, the denoisers in most iterative schemes (e.g., MMNet and OAMPNet) are specifically designed for i.i.d. Gaussian noise. Thus, ideally, the linear stage should avoid outputting correlated or non-Gaussian noise.

Fig. 5-1 shows the average noise power at the output of the linear and denoiser stages across iterations, for both OAMPNet and MMNet on 64×16 3GPP MIMO channels. The average noise power before and after the denoiser saturate at the same value in OAMPNet from the third layer ($t = 2$) onwards, showing that the denoisers are no longer effective after a few iterations.

We hypothesize that the reason OAMPNet’s denoisers become ineffective is that the noise distribution for OAMPNet is not Gaussian, whereas MMNet is able to provide near-Gaussian noise to its denoiser. We evaluate how close the noise distribution is to Gaussian for both schemes using the Anderson test [33]. In order to measure this score, we generate 10,000 samples per channel realization \mathbf{H} . We then calculate the Anderson score for the noise distribution at each linear stage output per transmitter, and per channel matrix. If this score is

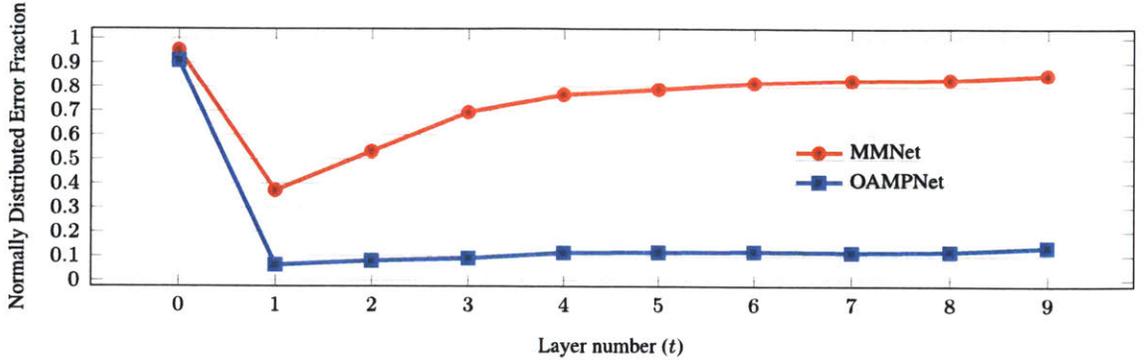


Figure 5-2: Percentage of transmitters that have Gaussian error distribution after the linear block for each layer with significance level of 5%. MMNet produces Normal-distributed errors at the output of linear blocks, while OAMPNet fails to achieve the Gaussian property.

below a threshold of 0.786, it indicates that the noise comes from a Gaussian distribution with a significance of 5%, i.e. the probability of false rejection of a Gaussian distribution is less than 5%.

In Fig. 5-2, we plot the average ratio of transmitters that have Normally-distributed noise at the output of the linear stage according to this test. Since in both schemes we start with $\hat{\mathbf{x}}_0 = 0$, the output of the linear stage at layer $t = 0$ is $\mathbf{A}_0 \mathbf{n}$, which is Gaussian. Thus, the fraction of transmitters with Gaussian noise is 1 in layer $t = 0$ for both schemes. However, both schemes deviate from Gaussian noise in layer $t = 1$ while sharply reducing the total noise power as seen in Fig. 5-1. However, MMNet deviates less from a Gaussian distribution. Unlike OAMPNet, in which the noise for 95% of the transmitters is not Gaussian at layer $t = 1$, for MMNet nearly 40% of the transmitter exhibit Gaussian noise. On the other hand, MMNet reduces the noise power slightly less than OAMPNet in layer $t = 1$.

In subsequent layers, the noise distribution for MMNet becomes increasingly Gaussian, with nearly 90% of transmitter passing the Anderson test by layer $t = 9$. By contrast, most transmitters in OAMPNet continue to exhibit non-Gaussian noise in subsequent layers, though the fraction of transmitter with Gaussian noise increases marginally.

Next, we measure the effect of input error power on the output error distribution of linear stages in both schemes. In other words, we are interested to know how $\|\mathbf{e}_{t-1}^{den}\|$ impacts the

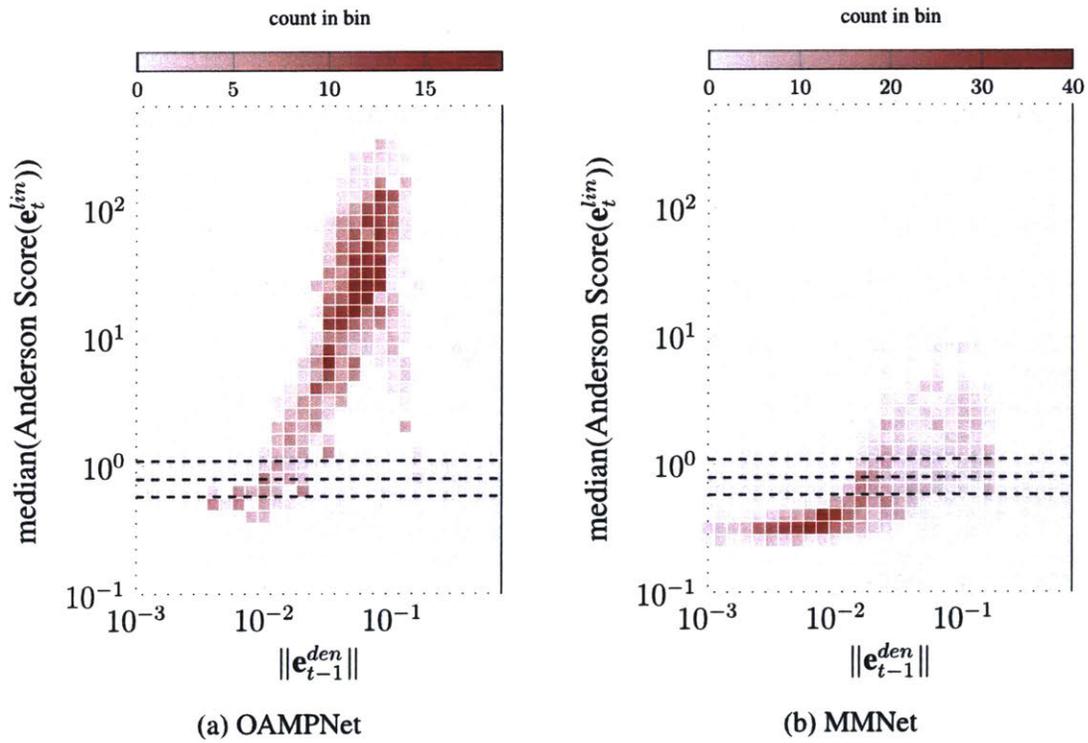


Figure 5-3: Median of Anderson score for the noise at the output of linear stage vs. the power of noise at the input of the stage. MMNet controls the linear block output error distribution to be Gaussian. Dashed horizontal lines show the thresholds for 1%, 5% and 15 significance level.

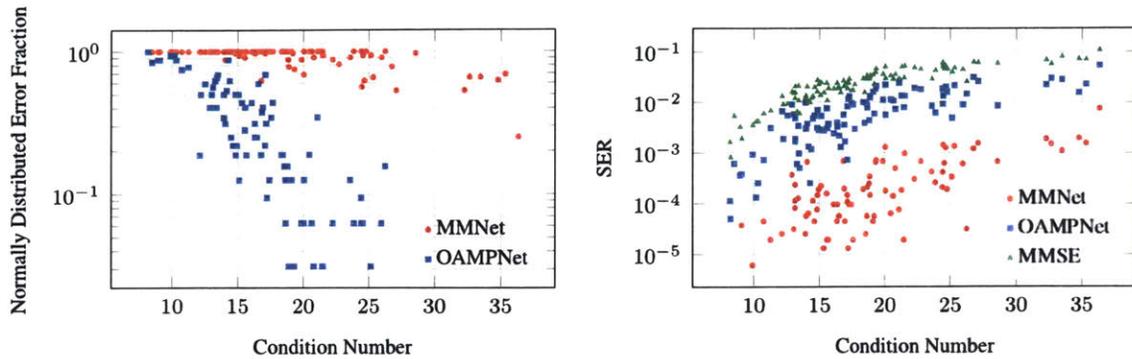
Gaussian distribution property of \mathbf{e}_t^{lin} . For this purpose, we choose the median of Anderson scores as a measure of the linear stage’s ability to control the distribution of its output errors. In Fig. 5-3, we show the 2D histogram of this median score for different values of $\|\mathbf{e}_{t-1}^{den}\|$. We also plot three thresholds of 1%, 5% and 15% significance for the normality test with dashed horizontal lines as a reference. To be Normally distributed with 1%, 5%, or 15% confidence, the Anderson scores must fall below the respective line.

We notice that the median score in both schemes increases with the norm of the error from the previous denoiser stage. In other words, the linear stages that have a higher input noise power produce outputs that deviate more from Gaussian noise. However, MMNet is $100\times$ better in terms of the median score at controlling the input error for large value of $\|\mathbf{e}_{t-1}^{den}\|$. This figure also suggests that the poor performance of OAMPNet in final layers is likely not only because of the aggressive approach it has taken at $t = 1$. Later linear stages are also not very good at controlling the distribution of their output errors.

5.3 Impact of channel condition number

Finally, we evaluate the impact of the channel *condition number* on MMNet and OAMPNet. A channel’s condition number is defined as the ratio of its largest singular value to the smallest. It is well-known to be more difficult to detect signals passing through channel matrices with higher condition numbers, as received signals belonging to different transmitted signals can be hard to distinguish passing through such channels.

OAMPNet tries to address this issue by introducing filtering over the singular values. In fact, the linear update equation in Eq. (2.10) attempts to map each singular value s in \mathbf{H} to $\theta_t^{(1)} s / (v_t^2 s^2 + \sigma^2)$. This in turn attempts to dynamically adjust the shape of the sphere mapping in each iteration by tuning $\theta_t^{(1)}$ and v_t^2 . We see that if all singular values are near each other, as is usually the case in i.i.d. Gaussian channels, OAMPNet easily maps each sphere of signals in the transmit space to a non-skewed sphere in the receive signals space. However, our results show that manipulating the singular values is not the best option for



(a) Ratio of e^{lin} Gaussians vs. channel condition number

(b) SER on 3GPP MIMO channels

Figure 5-4: Effect of condition number on performance of schemes. (a) MMNet is more robust in producing the right noise distribution for denoisers with changes in condition number. (b) SER is directly affected by the condition number.

poorer channel condition numbers.

In Fig. 5-4a, we show the scatterplot of performance of MMNet linear stage in terms of preserving the Gaussian distribution property of their output error distribution across different channel condition numbers sample from our 3GPP MIMO dataset after a few initial iterations ($t = 4$). We see that the fraction of Normally-distributed errors decreases quickly for OAMPNet with the increase in condition number, while MMNet maintains the ratio for a broader range of condition numbers in 3GPP MIMO channels. The consequence of this failure in meeting the underlying assumptions of the model shows up in Fig. 5-4b. In this figure, we show the scatterplot of SER at different condition numbers for 3GPP MIMO channels. Although all schemes' performances are affected by the condition number, MMNet can almost maintain an SER of around 10^{-3} almost across all conditions.

Chapter 6

Online Training Algorithm

Training deep-learning models is a computationally intensive task, often requiring hours or even days for large models. The computation overhead depends on two factors: (i) the total number of required training samples, and (ii) the size of the model. For example, in training a model like DetNet with order of 1M parameters, we need 50K iterations with a batch size of 5K samples, i.e., 250M training samples. If we assume each parameter of the model shows up in at least one floating-point operation per training sample, we require a minimum of 2.5×10^{14} floating-point operations for the entire process of training DetNet. This extreme computational complexity makes training such models online for each realization of \mathbf{H} all but impossible.

By contrast, MMNet has only $\sim 40\text{K}$ parameters, and training it from scratch requires about 1000 iterations with batch size of 500. Further, this chapter will show that by taking advantage of locality of the channels observed at a receiver, we can effectively train the model for each channel realization in 4 iterations (with batch size 500) on average. Thus training MMNet has 6 orders of magnitude lower computational overhead than DetNet, making online training for each realization of \mathbf{H} practical.

In this chapter, we first discuss the temporal and spectral locality of 3GPP MIMO channels. Next, we show how we can exploit these localities to accelerate online training.

6.1 Channel locality

The channel matrices measured at a base station exhibit two forms of locality:

- *Temporal*: Channel matrices change gradually over time as user devices move within a cell or the multipath environment changes. The samples of \mathbf{H} at nearby points in time are thus correlated.
- *Spectral*: A BS needs to recover signals from several frequency subcarriers (1024 in our 3GPP MIMO model). The channels for subcarriers in nearby frequencies are also strongly correlated, as the frequency merely affects the phase for multipath signal components incident on receiver antennas. For a path of length l , the phase difference for two subcarriers Δf apart in frequency is $l \cdot \Delta f$. Therefore the received signal and the channels for adjacent subcarriers will be similar at each receiver antenna.

Both forms of locality reduce the complexity of training for each channel realization, because (i) the cost of channel-specific computations can be amortized across multiple correlated channel realizations across time and frequency, (ii) the trained model for one channel realization can serve as strong initialization for training for adjacent channel realizations in the time-frequency plane.

Fig. 6-1 shows both forms of locality by plotting the correlation among the 3GPP MIMO channel samples across time and frequency (subcarriers) In order to compute these correlations, we have averaged the inner-product of the channel matrices with their shifted realization by the step value in the corresponding dimension. In this experiment, a shift of 1 step in the time domain corresponds to two channels at the same subcarrier frequency that are 118ms apart in time. Similarly, a shift of 1 step in the frequency domain corresponds to two channels at adjacent subcarriers (20MHz apart) at the same time. We normalize the inner-product by the norm of matrices, such that the correlation of a channel matrix with itself is 1. We observe a stronger locality in the frequency domain than in the time domain on these channels.

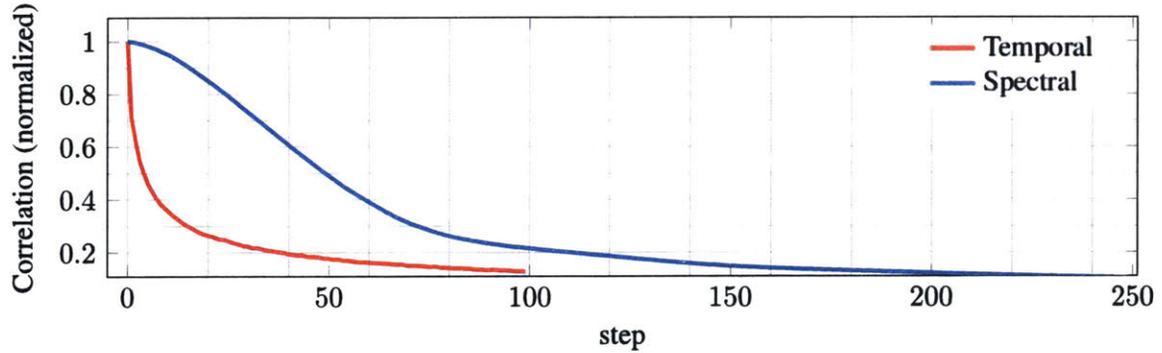


Figure 6-1: Correlation of channel samples across time and frequency dimensions. The correlation decays relatively quickly in the time dimension, but the channel matrices show strong locality across sub-carriers in frequency dimension.

6.2 Training algorithm and results

In this section, we show how channel locality can help reduce the total number of operations MMNet needs to decode each received signal \mathbf{y} at the BS. The computational complexity of MMNet is mostly dominated by the cost of online training for each new realization of the channel \mathbf{H} . This cost in turn depends on the channel *coherence time*. In the case of a quasi-static channel, as expected for instance in fixed-wireless access or backhaul solutions such as 5G home wireless (see Section 7.6.2 in [3]), the channel between the transmitter and receiver does not change for extended periods of time. In such cases, MMNet does not require frequent retraining and can reuse the same model until the communication channel changes significantly. However, MMNet can also operate at reasonable computation cost when the channel is changing fairly frequently. For example, our 3GPP MIMO channel samples were generated assuming all devices constantly move at a speed of 1 m/s, and after about 500ms, the channel correlation is less than 0.5. However, even in this scenario, we require only 4 training iterations on average per channel realization, as explained next.

To see how, note that a receiver at a BS must simultaneously decode signals from different subcarrier channels. Since channels exhibit strong correlations across sub-carriers (Fig. 6-1), training the MMNet detector on \mathbf{H} for one subcarrier produces a detector that will achieve very similar performance on adjacent subcarriers. The performance of this detector

Algorithm 1 MMNet online training

```
 $\mathcal{M} \leftarrow \text{model}(\text{MMNet})$  ▷ load MMNet model  
 $\mathcal{M}.\text{initialize}()$  ▷ Initialize model parameters randomly  
 $\text{batchsize} \leftarrow 500$  ▷ Training batch size  
 $n \leftarrow 0$  ▷  $n$  keeps time  
while True do  
  for  $f$  in  $\{1, 2, \dots, F\}$  do  
     $\mathbf{H}[f] \leftarrow$  Measured channel at time  $n$  and frequency  $f$   
     $\mathbf{D} \leftarrow$  Generate random  $(\mathbf{x}, \mathbf{y})$  pairs for  $\mathbf{H}[f]$  using Eq. (2.1)  
    if  $f == 1$  then  
       $\text{numTrainIterations} \leftarrow 1000$   
    else  
       $\text{numTrainIterations} \leftarrow 3$   
    end if  
     $\mathcal{M}.\text{train}(\mathbf{D}, \text{numTrainIterations}, \text{batchsize})$  ▷ Update model parameters  
     $\mathcal{M}_n[f] \leftarrow \mathcal{M}.\text{copy}()$   
  end for  
   $n \leftarrow n + 1$   
end while
```

will however decay for more distant subcarriers in the frequency domain.

Based on this observation, we propose the online training scheme in Algorithm 1. We start from a random initialization of the MMNet neural network model \mathcal{M} . We define n as an index for time intervals in which we can assume that channels do not change substantially. For each interval n , we measure a channel matrix $\mathbf{H}[f]$ for each subcarrier frequency f . The basic idea in the algorithm is to train the model for 1000 iterations (with a batch size of 500) for the first subcarrier ($f = 1$), then retrain the model using only 3 additional training iterations per subcarrier for all subsequent subcarriers. For each channel matrix $\mathbf{H}[f]$, we generate (\mathbf{x}, \mathbf{y}) training data pairs using Eq. (2.1). Once the model has trained for subcarrier f , we save a copy of the model as $\mathcal{M}_n[f]$ for detecting all signals received in time interval n on that subcarrier. We repeat the entire training algorithm in each time interval.

In Fig. 6-2, we show the result of this online training method on 3GPP MIMO channels for the QAM16 modulation. Roughly 95% of samples are decoded with SER of less than 0.02, while OAMPNet (trained explicitly for thousands of iterations on each realization of \mathbf{H} at each time interval n and subcarrier f) is only able to decode 50% of the samples at the same error rate. With this approach, MMNet performs 3.55K iterations of training

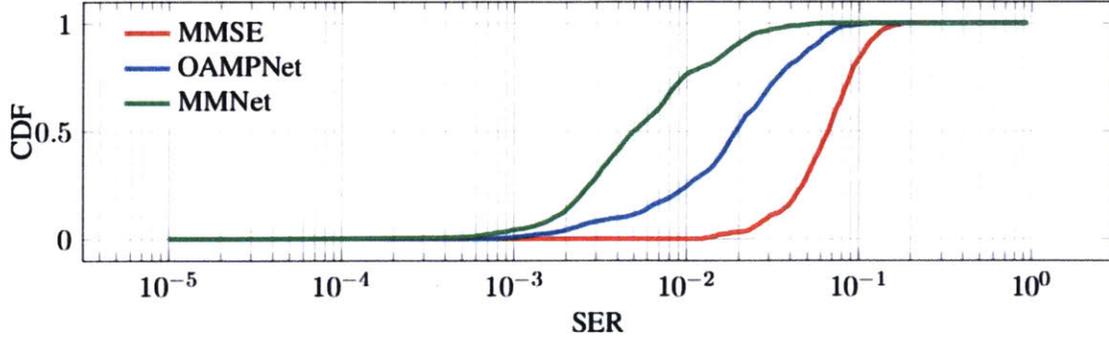


Figure 6-2: CDF of SER using Algorithm 1 for training MMNet on QAM16 modulation. MMNet requires only 4 overall iterations of batch size 500 per channel realization to train to a reasonable performance.

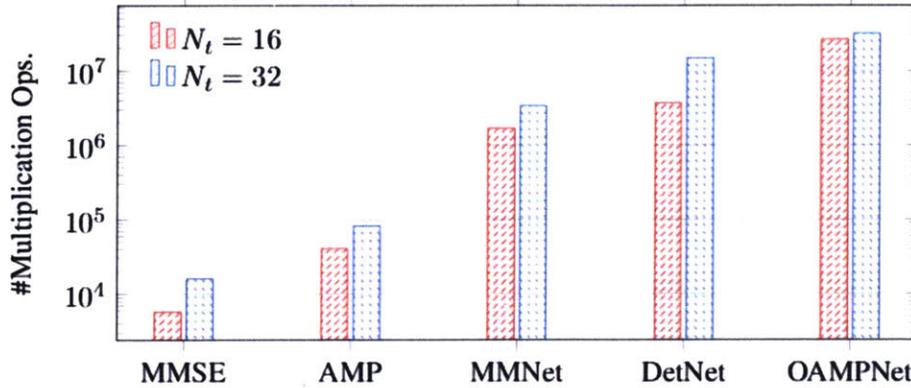


Figure 6-3: Number of multiplication operations per signal detection for different algorithms on QAM16 with $N_r = 64$ receive antennas in 3GPP MIMO model. Detection with MMNet, including its online training process, requires fewer multiplication operations than detection with pre-trained DetNet and OAMPNet models.

with batch size 500 in order to learn a detector for all 1024 subcarriers in total at each time interval n . Therefore the cost of online training is less than 4 iterations on average per channel realization, yet MMNet delivers better performance than other schemes, like OAMPNet and MMSE.

6.3 Computational complexity

One iteration of training MMNet on a batch of size of b has a complexity of $\mathcal{O}(bN_r^2)$, as detection takes $\mathcal{O}(N_r^2)$ in MMNet. To put this in perspective, a light-weight algorithm like

AMP has a complexity of $\mathcal{O}(N_r^2)$ dominated by the multiplication of the channel matrix and residual vectors. The MMSE scheme has a higher complexity of $\mathcal{O}(N_r^3)$ because it needs to invert a matrix. OAMPNet similarly requires a matrix inversion, resulting in a complexity of $\mathcal{O}(N_r^3)$.

Moving beyond $\mathcal{O}(\cdot)$ analysis, Fig. 6-3 shows the average number of multiplication operations required per signal detection on 3GPP MIMO channels for learning-based algorithms in addition to two classic baselines, MMSE and AMP. All algorithms may reuse their computation if feasible. In particular, in every interval of channel coherence in 3GPP MIMO model, each algorithm receives ~ 100 signals to detect [3, Definition 2.2 on page 220]. MMSE takes the best advantage of this as it calculates the required channel matrix inversion only once for all 100 received signals in the coherence interval. This way, MMSE depreciates its principal bottleneck calculation by a factor of $100\times$ resulting in $5\text{-}7\times$ less computation than AMP algorithm which cannot reuse the calculations but still has very moderate computation overhead by design. MMNet with online training and detection operations together, places after AMP with $2\text{-}5\times$ fewer multiplications than pre-trained DetNet. However, neither AMP nor DetNet extend to spatially correlated channels. MMNet reuses the weights it trains with 4 iterations of batch size 500 for all 100 received signals in a coherence interval. However, unlike MMSE, OAMPNet has to calculate a new matrix inversion in each layer for every received signal as v_t^2 in Eq. (2.10) depends on $\hat{\mathbf{x}}_t$.

Consequently, the cost of MMNet with its online training algorithm is $10\text{-}15\times$ less than OAMPNet depending on the system size. MMNet has only $41\times$ higher computational complexity than a very light iterative approach like AMP, which only works near Maximum-Likelihood under specific circumstances of i.i.d. Gaussian channels.

Chapter 7

Conclusion

This dissertation proposed a new deep learning architecture for massive MIMO detection with an online training algorithm. MMNet outperforms state-of-the-art detection algorithms on a wide combination of channel conditions, modulations, and system size. We designed MMNet as an iterative algorithm and showed that the right degree of flexibility in the model in addition to leveraging the channel's locality, can enable the idea of online training at a less or equal computation complexity of other deep-learning based schemes like OAMPNet. MMNet is 4-8dB better overall than the classic MMSE detector and it requires 2.5dB lower SNR at the same SER, relative to the second-best detection scheme, OAMPNet, at 10-15 \times less computational complexity.

Bibliography

- [1] Hengtao He, Chao-Kai Wen, Shi Jin, and Geoffrey Ye Li. A model-driven deep learning network for mimo detection. *arXiv preprint arXiv:1809.09336*, 2018.
- [2] Afif Osseiran, Jose F Monserrat, and Patrick Marsch. *5G Mobile and Wireless Communications Technology*. Cambridge University Press, 2016.
- [3] Emil Björnson, Jakob Hoydis, and Luca Sanguinetti. Massive MIMO networks: Spectral, energy, and hardware efficiency. *Foundations and Trends® in Signal Processing*, 11(3-4):154–655, 2017. ISSN 1932-8346. doi: 10.1561/20000000093. URL <http://dx.doi.org/10.1561/20000000093>.
- [4] Emre Telatar. Capacity of Multi-antenna Gaussian Channels. *Transactions on Emerging Telecommunications Technologies*, 10(6):585–595, 1999.
- [5] Emanuele Viterbo and Joseph Boutros. A universal lattice code decoder for fading channels. *IEEE Transactions on Information theory*, 45(5):1639–1642, 1999.
- [6] Shaoshi Yang and Lajos Hanzo. Fifty years of mimo detection: The road to large-scale mimos. *IEEE Communications Surveys & Tutorials*, 17(4):1941–1988, 2015.
- [7] Erik G Larsson. MIMO detection methods: How they work. *IEEE signal processing magazine*, 26(3), 2009.
- [8] Ami Wiesel, Yonina C Eldar, and Shlomo Shamai. Semidefinite relaxation for detection of 16-qam signaling in mimo channels. *IEEE Signal Processing Letters*, 12(9): 653–656, 2005.

- [9] Neev Samuel, Tzvi Diskin, and Ami Wiesel. Deep mimo detection. *arXiv preprint arXiv:1706.01151*, 2017.
- [10] Junjie Ma and Li Ping. Orthogonal amp. *IEEE Access*, 5:2020–2033, 2017.
- [11] Sergey L Loyka. Channel capacity of mimo architecture using the exponential correlation matrix. *IEEE Communications letters*, 5(9):369–371, 2001.
- [12] 3GPP TR 36.873. 2015. “Study on 3D channel model for LTE”. Technical report.
- [13] Stephan Jaeckel, Leszek Raschkowski, Kai Börner, and Lars Thiele. QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials. *IEEE Trans. Antennas Propag.*, 62(6):3242–3256, 2014.
- [14] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- [15] Charles Jeon, Ramina Ghods, Arian Maleki, and Christoph Studer. Optimality of large mimo detection via approximate message passing. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 1227–1231. IEEE, 2015.
- [16] Lajos Hanzo, Soon Xin Ng, WT Webb, and T Keller. *Quadrature amplitude modulation: From basics to adaptive trellis-coded, turbo-equalised and space-time coded OFDM, CDMA and MC-CDMA systems*. IEEE Press-John Wiley, 2004.
- [17] Alberto Del Pia, Santanu S Dey, and Marco Molinaro. Mixed-integer quadratic programming is in np. *Mathematical Programming*, 162(1-2):225–240, 2017.
- [18] Amine Mezghani and Josef A Nossek. Belief propagation based mimo detection operating on quantized channel output. In *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pages 2113–2117. IEEE, 2010.
- [19] Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.

- [20] Antonia M Tulino, Sergio Verdú, et al. Random matrix theory and wireless communications. *Foundations and Trends® in Communications and Information Theory*, 1 (1):1–182, 2004.
- [21] Peter W Wolniansky, Gerard J Foschini, GD Golden, and Reinaldo A Valenzuela. V-blast: An architecture for realizing very high data rates over the rich-scattering wireless channel. In *Signals, Systems, and Electronics, 1998. ISSSE 98. 1998 URSI International Symposium on*, pages 295–300. IEEE, 1998.
- [22] Mahesh K Varanasi and Behnaam Aazhang. Multistage detection in asynchronous code-division multiple-access communications. *IEEE Transactions on communications*, 38(4):509–519, 1990.
- [23] WH Chin, AG Constantinides, and DB Ward. Parallel multistage detection for multiple antenna wireless systems. *Electronics Letters*, 38(12):597–599, 2002.
- [24] Ori Shental, Sivarama Venkatesan, Alexei Ashikhmin, and Reinaldo A Valenzuela. Massive blast: An architecture for realizing ultra-high data rates for large-scale mimo. *arXiv preprint arXiv:1708.05405*, 2017.
- [25] Toshiyuki Tanaka and Masato Okada. Approximate belief propagation, density evolution, and statistical neurodynamics for cdma multiuser detection. *IEEE Transactions on Information Theory*, 51(2):700–706, 2005.
- [26] Neev Samuel, Tzvi Diskin, and Ami Wiesel. Learning to detect. *arXiv preprint arXiv:1805.07631*, 2018.
- [27] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.
- [28] Wing-Kin Ma, Chao-Cheng Su, Joakim Jaldén, and Chong-Yung Chi. Some results on 16-qam mimo detection using semidefinite relaxation. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2673–2676. IEEE, 2008.

- [29] Inc. Gurobi Optimization. Gurobi optimizer reference manual, 2016. URL <http://www.gurobi.com>.
- [30] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- [31] Wing-Kin Ma, Chao-Cheng Su, Joakim Jaldén, Tsung-Hui Chang, and Chong-Yung Chi. The equivalence of semidefinite relaxation mimo detectors for higher-order qam. *IEEE Journal of Selected Topics in Signal Processing*, 3(6):1038–1052, 2009.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Theodore W Anderson, Donald A Darling, et al. Asymptotic theory of certain "goodness of fit" criteria based on stochastic processes. *The annals of mathematical statistics*, 23(2):193–212, 1952.