



CHALMERS
UNIVERSITY OF TECHNOLOGY



Sparse sampling for massive MIMO downlink CSI reference signals

Submodular optimisation algorithms for near-optimal sparse
signal design and neural network-based channel reconstruction

Master's thesis in Complex adaptive systems

ANDREAS FÜHR

DEPARTMENT OF ELECTRICAL ENGINEERING

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

MASTER'S THESIS 2025

Sparse sampling for massive MIMO downlink CSI reference signals

Submodular optimisation algorithms for near-optimal sparse
signal design and neural network-based channel reconstruction

ANDREAS FÜHR



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Communication, Antennas and Optical Networks
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025

Sparse sampling for massive MIMO downlink CSI reference signals

Submodular optimisation algorithms for near-optimal sparse signal design and neural network-based channel reconstruction

ANDREAS FÜHR

© 2025 ANDREAS FÜHR. All rights reserved.

Supervisors: Dr. Xinlin Zhang, Dr. Johan Wings, and Dr. Keerthi Kumar Nagalapur,
Ericsson AB

Advisor: Kaan Okumus, Department of Electrical Engineering

Examiner: Prof. Erik Ström, Department of Electrical Engineering

Master's thesis 2025

Department of Electrical Engineering

Division of Communication, Antennas and Optical Networks

Chalmers University of Technology

SE-412 96 Gothenburg

Telephone +46 31 772 1000

Typeset in L^AT_EX

Gothenburg, Sweden 2025

Sparse sampling for massive MIMO downlink CSI reference signals

Submodular optimisation algorithms for near-optimal sparse signal design and neural network-based reconstruction

ANDREAS FÜHR

Department of Electrical Engineering
Chalmers University of Technology

Abstract

The number of antennas used in massive multiple input multiple output (MIMO) systems is expected to increase significantly to meet requirements of future radio access networks (RANs). In legacy 5G New Radio, scaling the number of antennas imposes a proportional increase in overhead associated with the acquisition of channel state information (CSI) through downlink (DL) transmission of reference signals (CSI-RS). This thesis investigates methods to reduce the overhead of CSI-RS transmission using a twofold approach: optimising pilot placement for sparse sounding of the reference signals, and reconstructing the full channel information from these sparse measurements at the user equipment (UE).

First, the sparse sampling of CSI-RS is formulated as a submodular optimisation problem, and a cost function is presented based on the frame potential of the DL channel estimated by the UE. A greedy algorithm for solving the sparse pilot placement problem is proposed and evaluated for a simulated 3rd Generation Partnership Project (3GPP) MIMO urban microcell environment with a uniform planar array (UPA), achieving near-optimal pilot placement for subsets of antenna ports.

The second part of the thesis investigates the recovery of the full channel information from the sparsely sounded CSI-RS using an artificial neural network (ANN). A physics-informed U-Net architecture is developed, that leverages the sparse angular representation of the DL channel to recover the full-rank channel. The ANN is trained on a large dataset of simulated noiseless DL channels for the same 3GPP environment and for several different spatial pilot configurations and muting levels.

The results of the experiments show that the ANN model can achieve a reconstruction accuracy comparable to basis pursuit denoising (BPDN), while outperforming BPDN in computational efficiency. In addition, the choice of spatial antenna port muting pattern has a noticeable impact on the reconstruction performance of both methods in the considered scenario, with the found near-optimal sampling patterns gives the closest spectral similarity to the full-rank channel in terms of the Itakura-Saito distance. The combined approach of optimising sparse pilot placement and using a neural network for reconstruction demonstrates the potential of AI functionality for CSI-RS overhead reduction and for improving the performance of massive MIMO systems in upcoming 6G networks and beyond.

Keywords: massive MIMO, CSI, sparse sampling, antenna selection, submodular optimisation, frame potential, channel estimation, neural networks, U-Net, basis pursuit denoising, machine learning, 6G

Acknowledgements

I would like to express my sincere gratitude to Ericsson Research in Gothenburg for the opportunity to work on this thesis project as part of the Ericsson Master Thesis Scholarship program, and in particular to research manager Henrik Sahlin for arranging the thesis project in his team on an interesting and challenging topic. My supervisors at Ericsson Research, Xinlin Zhang, Johan Wings, and Keerthi Kumar Nagalapur, have provided very valuable support, encouragement and guidance throughout the project. Their expertise and insights have enriched my understanding of the subject matter and has been instrumental in shaping the direction of the work. I am also grateful to my advisor Kaan Okumus at the Department of Electrical Engineering for his assistance and support during the project, especially for fruitful discussions and helpful inputs on the sparse sampling part of the project. To my examiner at the Department of Electrical Engineering, Erik Ström, I would like to extend my appreciation for his thoughtful feedback and encouragement throughout the project. Lastly, I would like to acknowledge Vinnova for funding the supervision of this thesis project.

Andreas Führ
Gothenburg, May 2025

List of acronyms

A list of acronyms that have been used throughout this work is given in alphabetical order.

2D, 3D	[Two, Three]-dimensional.
3G, 4G, 5G, 6G	[Third, Fourth, Fifth, Sixth] generation (mobile network technology).
AI	Artificial intelligence.
ANN	Artificial neural network.
AoA	Angle of arrival.
AoD	Angle of departure.
AoSA	Array of subarrays.
BP	Basis pursuit.
BPDN	Basis pursuit denoising.
BS	Base station.
CNN	Convolutional neural network.
CS	Compressed sensing.
CSI	Channel state information.
CSI-RS	CSI reference signal.
DFT	Discrete Fourier transform.
DL	Downlink.
FDD	Frequency division duplexing.
FP	Frame potential.
FUNTF	Finite unit-norm tight frame.
LP	Linear program.
LOS	Line-of-sight.
MIMO	Multiple input multiple output.
Min-RVLS	Minimum relevant variables in linear system.
ML	Machine learning.
MSE	Mean squared error.

MU-MIMO	Multi-user multiple input multiple output.
NMSE	Normalised mean squared error.
NLOS	Non-line-of-sight.
NR	New radio.
OFDM	Orthogonal frequency division multiplexing.
OH	Overhead.
OMP	Orthogonal matching pursuit.
ONB	Orthonormal basis.
PINN	Physics-informed (artificial) neural network.
PRB	Physical resource block.
PSD	Power spectral density.
RAN	Radio access network.
ReLU	Rectified linear unit.
RF	Radio frequency.
SA	Subarray.
SINR	Signal-to-interference-and-noise ratio
SNR	Signal-to-noise ratio.
SU-MIMO	Single-user multiple input multiple output.
TDD	Time division duplexing.
UE	User equipment.
UL	Uplink.
UMi	Urban microcell.
UPA	Uniform planar array.

Nomenclature

Sets, operators, quantities and parameters that have been used throughout this work are listed below.

Sets and spaces

\mathbb{R}	Real numbers.
\mathbb{C}^N	Complex vector space of dimension N .
\mathcal{H}	Hilbert space
\mathcal{N}	Set of indices for the source signal.
\mathcal{L}	Sampling set of the source signal, $\subseteq \mathcal{N}$.
\mathcal{S}	$= \mathcal{N} \setminus \mathcal{L}$. Complement of \mathcal{L} .

Indices, dimensions and variables

i, j, k, l, o, s	Element indices.
K, L, M, N, P, Q, S	Dimensions of a tensor, or number of elements in a set.
\mathbf{x}	Vector in \mathbb{C}^N .
x_i	i th element of \mathbf{x} .
\mathbf{A}	Matrix in $\mathbb{C}^{N \times M}$, or tensor in $\mathbb{C}^{N_1 \times \dots \times N_n}$.
A_{ij}	(i, j) th element of a matrix \mathbf{A} .
\mathbf{a}_i	i th column of a matrix \mathbf{A} .
$\mathbf{A}^{(k)}$	k th realisation of \mathbf{A} over a dataset.
$[\mathbf{A}]_{ij}$	(i, j) th block of a block matrix \mathbf{A} .
j	Imaginary unit, $j \equiv \sqrt{-1}$.
c	Speed of light.
\mathbf{I}_N	Identity matrix of size $N \times N$.
$\mathbf{1}_N$	Vector of ones of size N .

Operators

$ \cdot $	Absolute value of a complex number, or cardinality of a set.
$\ \cdot\ _p$	ℓ_p -norm of a vector.
\otimes	Kronecker product.
\oplus	Matrix direct sum.
\odot	Hadamard product.
\circ	Hadamard power.
$\text{vec}(\cdot)$	Vectorisation.
$\text{vec}_{(N_i, \dots, N_j)}(\cdot)$	Vectorisation.
$\text{vec}_{(N_1, \dots, N_n)}(\cdot)$	Partial vectorisation over dimensions N_1, \dots, N_n .
$\text{diag}(\cdot)$	Diagonal matrix.
$\text{tr}(\cdot)$	Trace.
\mathbf{A}^\top	Transpose of \mathbf{A} .
\mathbf{A}^H	Hermitian transpose of \mathbf{A} .
$\mathcal{F}\{\cdot\}$	Discrete Fourier transform.
$\mathbb{E}[\cdot]$	Arithmetic mean.
$\mathbb{E}_w[\cdot]$	Weighted mean with weights w .
$\text{Cov}(\cdot, \cdot)$	Covariance.
$\text{FP}_p(\cdot)$	p -frame potential.
$\Re(\cdot)$	Real part of a complex number.
$\Im(\cdot)$	Imaginary part of a complex number.
$\langle \cdot, \cdot \rangle$	Inner product of two vectors.

Special parameters

$\varphi_{\mathcal{L}}$	Sampling vector for a sampling set \mathcal{L} .
$\Phi_{\mathcal{L}}$	Sampling matrix for a sampling set \mathcal{L} .
\mathbf{H}	Downlink MIMO channel.
\mathbf{H}_{bs}	Downlink MIMO channel in beamspace.

Contents

List of acronyms	ix
Nomenclature	x
Contents	xiii
List of figures	xv
List of tables	xvi
List of algorithms	xvii
1 Introduction	1
1.1 Objective	2
1.2 Limitations	2
1.3 Outline	3
2 Background	5
2.1 Mathematical preliminaries	5
2.1.1 Direct sum of matrices and the Kronecker product	5
2.1.2 Tensor vectorisation	6
2.1.3 The multidimensional discrete Fourier transform on matrix-vector form	7
2.2 Massive MIMO and CSI acquisition	8
2.3 System model	9
3 Sparse sampling of CSI-RS using submodular optimisation	15
3.1 Sparse sampling for linear inverse problems	15
3.2 Submodular optimisation for designing near-optimal sampling schemes	17
3.3 The generalised average-case criterion for stochastic systems	19
3.4 The p -frame potential	20
3.5 Related work	22
3.6 An objective function based on the p -frame potential	23
3.7 Sparse sampling of the DL channel	26
4 Neural network-based reconstruction of sparsely sounded CSI-RS	29

4.1	Reconstruction of sparsely sampled linear inverse problems	29
4.2	Classical compressed sensing-based reconstruction algorithms	30
4.3	Compressed sensing-based reconstruction of sparsely sampled DL channel	32
4.4	A distance measure for the reconstruction problem with sparsity in the angular domain	32
4.5	A PINN reconstruction model for DL channel reconstruction	33
5	Experiments and results	37
5.1	Simulation of DL channels and data partitioning	37
5.2	Fixed spatial sampling patterns for CSI-RS	39
5.3	Near-optimal spatial sampling patterns for CSI-RS	41
5.4	Neural network-based reconstruction of the DL channel for CSI-RS	46
5.4.1	Pre-training and fine-tuning of the neural network	46
5.4.2	Reconstruction of DL channels	46
6	Conclusion	53
6.1	Future work	54
	Bibliography	55

List of figures

2.1	Antenna configuration of a massive MIMO system at link level	10
2.2	Illustration of a $(4 \times 1)_{SA}(2 \times 8)$ UPA, divided into SAs with one RF chain per polarisation and SA.	12
3.1	A sampling scheme with a sampling set \mathcal{L} and the corresponding sampling vector $\boldsymbol{\varphi}_{\mathcal{L}}$	16
3.2	A sampling scheme for the sparsely sampled linear inverse problem with a sampling set \mathcal{L} and corresponding sampling matrix $\boldsymbol{\Phi}_{\mathcal{L}}$	17
4.1	Half-sphere illustration of the DL channel in TX beamspace	34
4.2	ReconstructionUNet architecture overview for the sparse DL channel estimation.	35
4.3	Convolutional block (ConvBlock) of the U-Net architecture.	35
5.1	Hexagonal cell layout for BS sites	37
5.2	TX radiation patterns for the UMi simulation	39
5.3	Data partitioning for simulated DL channels	39
5.4	Randomised spatial antenna port sampling patterns	40
5.5	Randomised spatial antenna port sampling patterns jointly across polarisations	40
5.6	Checked spatial antenna port sampling patterns	40
5.7	Computed submodular function values for Algorithm 3.3 with upper and lower bounds across the test dataset, with an $N_t = 8 \times 8 \times 2 = 128$ port UPA and separate muting for polarisations.	41
5.8	Near-optimal spatial antenna port sampling patterns from submodular optimisation using Algorithm 3.3	42
5.9	Itakura-Saito distance for the near-optimal muting patterns with 50 % muting level, as compared to the fixed sampling patterns.	43
5.10	Itakura-Saito distance for the near-optimal muting patterns with 25 % muting level, as compared to the fixed random sampling patterns.	43
5.11	TX radiation patterns for the fixed UPA sampling patterns for the 25 % and 50 % muting levels.	44
5.12	Reconstruction performance of ReconstructionUNet and BPDN (Itakura-Saito distance), 50 % muting level	48

5.13 Reconstruction performance of ReconstructionUNet and BPDN (Itakura-Saito distance), 25 % muting level	48
5.14 TX radiation patterns for the reconstruction of a DL channel realisation with fixed UPA sampling patterns for 25 % and 50 % muting levels using ReconstructionUNet and BPDN.	49

List of tables

5.1	Simulation parameters for the UMi simulation scenario.	38
5.2	Average computational performance of ReconstructionUNet and BPDN per downlink channel.	47
5.3	Reconstruction performance of ReconstructionUNet and BPDN (NMSE) .	47

List of algorithms

3.1	Greedy algorithm for submodular function maximisation.	19
3.2	Greedy algorithm for the sampling scheme.	25
3.3	Greedy algorithm for the generalised average-case criterion sampling scheme.	26

1

Introduction

By the end of 2024, the total monthly global mobile network data traffic exceeded 157 exabytes (EB) and it is predicted to increase to 303 EB per month by 2030 [1]. Massive multiple input multiple output (MIMO) technology, utilising a large number of transmit (TX) and/or receive (RX) antennas, is a centrepiece for developing advanced antenna systems that can meet the requirements of future networks. One factor that determines the growth potential of mobile networks is the performance of the deployed networks, which can be improved by increasing the spatial and spectral efficiencies of the networks. In radio access networks (RANs) using multi-antenna technology, knowledge about the properties of the propagation channel is essential for enabling advanced beamforming and scheduling schemes [2, Ch. 6.4]. In legacy 5G radio access called New Radio (NR), developed by the 3rd Generation Partnership Project (3GPP), a base station (BS) can obtain so-called channel state information (CSI) via feedback from the user equipment (UE). In such an operation, termed closed-loop CSI feedback, the UE measures a number of CSI reference signals (CSI-RSs) transmitted from antennas at the BS, from which a CSI report is derived and transmitted to the BS. In feedback-based CSI acquisition, the number of reference signals required increases proportionally with the number of antennas at the BS. Due to the benefits of massive MIMO technology, it is expected that antennas with significantly more antenna elements will be deployed at both the BS and UE in future networks; wide antenna arrays at the BS with many columns may for example increase performance in networks where UEs are spread out horizontally [2, Ch. 15.2]. As a result, large TX antenna arrays may introduce significant overhead for acquisition of CSI, thereby leaving fewer resources for data transmission. To reduce the overhead of CSI-RSs in massive MIMO scenarios, antenna ports may be sounded in a sparse manner, i.e., only a spatial subset of the total number of antenna ports is used as pilot signals and the remaining ports are muted, for example according to the method proposed in [3]. The task at hand is then to reconstruct the full DL channel at the UE side estimated partial DL channel from the measured sparse CSI-RSs.

Starting with the 3GPP Release 18 with the introduction of 5G-Advanced in NR [4], there have been ongoing efforts to study artificial intelligence (AI) and ML techniques for overhead reduction in CSI reporting, as one of three selected AI/ML use cases for the radio interface that is expected to be foundational for future generations of RANs [5]. In this context, deep learning methods for UE-side estimation of the channel matrix is an important area of research for gaining an understanding of the possibility of AI models to reduce overhead in CSI feedback. However, the dual compression-decompression problem of sparse pilot placement and reconstructing the full DL channel from sparse

1. Introduction

measurements has not been extensively studied in the literature, in particular regarding AI-based algorithms. Secondly, finding optimal pilot placement schemes for massive MIMO has not been studied in detail, and existing work does not emphasise or consider the use of sparsely sounded CSI-RSs. Additionally, there has been limited work on developing artificial neural networks (ANNs) for channel reconstruction from a physics-informed perspective, so-called physics-informed neural networks (PINNs), where prior knowledge about the properties of the physical system is incorporated into the model architecture. The contributions of this master thesis consists of addressing these research gaps and demonstrating the potential of the developed methods for overhead reduction in massive MIMO systems.

1.1 Objective

The aim of this master thesis project has been to investigate methods for reducing the overhead of CSI-RSs in massive MIMO systems using a twofold approach: optimising pilot placement for sparse sounding of the reference signals, and reconstructing the full channel information from these sparse measurements at the UE side. To achieve this, we have asked the following research questions: How can sparse pilot placement be optimised for massive MIMO systems? How can the full DL channel be reconstructed from the sparse measurements at the UE side using deep learning methods, and what are suitable neural network architectures for this task? How do different pilot selection strategies affect the performance of the models? What are the performance trade-offs of the proposed methods compared to classical ML methods based on compressed sensing (CS)? What are suitable measures of reconstruction performance for the problem of reconstructing the full DL channel from sparse measurements? What are the implications of the results for reducing the DL reference signal overhead for CSI acquisition in massive MIMO systems?

1.2 Limitations

The scope of this thesis project is limited in a number of aspects. We have only considered feedback-based CSI acquisition, i.e., the BS transmits CSI-RSs to the UE and the UE produces and sends back a CSI report to the BS for TX precoding. Furthermore, we assumed ideal, noiseless conditions for the measurement of CSI-RSs, i.e., perfect CSI at the UE. The two parts of this thesis project, sparse pilot placement and reconstruction of the full DL channel, are evaluated separately and for one standardised MIMO scenario. Also, temporal correlations and time-varying properties of the channel matrix are disregarded, i.e. the realisations of the channel matrix for a single instant in time are used for training and evaluation. This means that the problem of channel aging and *prediction* of future CSI from historical data is not addressed. Lastly, to reduce the scope of the thesis project, only one type of AI model and CS-based benchmark method is considered.

1.3 Outline

This thesis report is structured as follows: We first introduce some mathematical notation used throughout the work, give a brief background to the topic of massive MIMO and CSI acquisition, after which a system model for CSI-RS transmission is given. We then present the two main parts of the thesis in Chapters 3 and 4: The first chapter focuses on the sparse pilot placement problem, where we formulate the problem as a general sparse sampling problem for a linear multidimensional map representing a physical system. We derive a cost function based on the frame potential and propose a greedy algorithm for solving the sparse pilot placement problem given a dataset of realisations of the physical system. The sparse sampling theory is then applied to the problem of optimising the placement of CSI-RS in massive MIMO systems. Chapter 4 presents the second part of the thesis, where we investigate the recovery of the sparsely sampled physical system, with direct application to the DL channel reconstruction from partial CSI-RS sounding. We give an introduction to classical reconstruction methods based on CS and apply the basis pursuit denoising (BPDN) algorithm to the problem of reconstructing the full DL channel from a sparse sampling scheme of the pilot signals. We then propose a physics-informed U-Net architecture that leverages the sparse angular representation of the DL channel to recover the full channel information.

To validate the proposed methods, we conducted a series of experiments presented in Chapter 5. DL channels were simulated for a 3GPP urban microcell (UMi) environment with a uniform planar array (UPA) at the BS and a single antenna at the UE. The near-optimal sparse antenna port sampling pattern found by the proposed greedy algorithm for sparse pilot placement, evaluated for a number of different pilot placements and muting levels, is compared to a number of fixed pilot placement designs. For the ANN-based reconstruction, we trained the U-Net on a large dataset of simulated noiseless DL channels for the same 3GPP environment and for several different spatial pilot configurations and muting levels. The performance of the ANN model is evaluated in terms of reconstruction accuracy and computational efficiency, and compared to the performance of BPDN.

1. Introduction

2

Background

2.1 Mathematical preliminaries

The properly structured sparse sampling of and linear transforms on the multidimensional DL channel presented in this thesis necessitate a number of non-trivial mathematical preliminaries. The operators presented in the following sections are the matrix direct sum, the Kronecker product, and the vectorisation of tensors. Furthermore, using these operators, we introduce a matrix-vector formulation of general multilinear maps, with a specific application to the n -dimensional discrete Fourier transform (DFT) and its inverse.

2.1.1 Direct sum of matrices and the Kronecker product

Given two matrices $\mathbf{A} \in \mathbb{C}^{N_1 \times M_1}$ and $\mathbf{B} \in \mathbb{C}^{N_2 \times M_2}$, the *matrix direct sum* is a block diagonal matrix that combines the two matrices into a larger matrix according to

$$\mathbf{A} \oplus \mathbf{B} := \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{bmatrix} \in \mathbb{C}^{(N_1+N_2) \times (M_1+M_2)}, \quad (2.1)$$

where the $\mathbf{0}$'s are block zero matrices of appropriate dimensions. The *Kronecker product* of two matrices $\mathbf{A} \in \mathbb{C}^{N_1 \times M_1}$ and $\mathbf{B} \in \mathbb{C}^{N_2 \times M_2}$ is defined as the block matrix

$$\mathbf{A} \otimes \mathbf{B} := \begin{bmatrix} A_{11}\mathbf{B} & \cdots & A_{1M_1}\mathbf{B} \\ \vdots & \ddots & \vdots \\ A_{N_11}\mathbf{B} & \cdots & A_{N_1M_1}\mathbf{B} \end{bmatrix} \in \mathbb{C}^{N_2 N_1 \times M_2 M_1}, \quad (2.2)$$

where A_{ij} is the (i, j) th element of \mathbf{A} , for $i = 1, \dots, N_1$ and $j = 1, \dots, M_1$. Properties of the Kronecker product that we will utilise in this work are associativity, distributivity over transpose and conjugate transpose, invertibility, and constant factorisation, respectively given by

$$(\mathbf{A} \otimes \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes (\mathbf{B} \otimes \mathbf{C}), \quad (2.3a)$$

$$(\mathbf{A} \otimes \mathbf{B})^\top = \mathbf{A}^\top \otimes \mathbf{B}^\top, \quad (2.3b)$$

$$(\mathbf{A} \otimes \mathbf{B})^H = \mathbf{A}^H \otimes \mathbf{B}^H, \quad (2.3c)$$

$$(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1}, \quad (2.3d)$$

$$(c\mathbf{A}) \otimes \mathbf{B} = (\mathbf{A} \otimes c\mathbf{B}) = c(\mathbf{A} \otimes \mathbf{B}). \quad (2.3e)$$

2. Background

2.1.2 Tensor vectorisation

The vectorisation of a matrix is a common operation in linear algebra. To generalise this operation to higher-order tensors, we here define a partial vectorisation operator, and its inverse operation. The *vectorisation* $\text{vec} : \mathbb{C}^{N_1 \times N_2} \rightarrow \mathbb{C}^{N_1 N_2}$ of a matrix $\mathbf{A} \in \mathbb{C}^{N_1 \times N_2}$ is defined as the column-wise stacking of the columns of \mathbf{A} ;

$$\text{vec}(\mathbf{A}) := \begin{bmatrix} \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_{N_2} \end{bmatrix} \in \mathbb{C}^{N_1 N_2}, \quad (2.4)$$

where $\mathbf{a}_j \in \mathbb{C}^{N_1}$ is the j th column of \mathbf{A} , for $j = 1, \dots, N_2$. The inverse operation of vectorisation, termed *matrixisation*, reshapes a vector into a matrix, defined as

$$\text{vec}_{\{1,2\}}^{-1}(\mathbf{b}) := \mathbf{B} \in \mathbb{C}^{N_1 \times N_2}, \quad (2.5)$$

where $\mathbf{b} \in \mathbb{C}^{N_1 N_2}$ is a vector, $\{1, 2\}$ indexes the dimensions N_1, N_2 , and $\mathbf{B} \in \mathbb{C}^{N_1 \times N_2}$ is the matrix obtained by unstacking N_1 consecutive elements of \mathbf{b} into N_1 columns of \mathbf{B} , for $N_2 = \dim(\mathbf{b})/N_1$. These operations satisfy

$$\text{vec}_{\{1,2\}}^{-1}(\text{vec}(\mathbf{A})) = \mathbf{A}, \quad (2.6)$$

$$\text{vec}\left(\text{vec}_{\{1,2\}}^{-1}(\mathbf{b})\right) = \mathbf{b}, \quad (2.7)$$

and scale linearly:

$$\text{vec}(c\mathbf{A}) = c \cdot \text{vec}(\mathbf{A}), \quad \text{vec}_{\{1,2\}}^{-1}(c\mathbf{b}) = c \cdot \text{vec}_{\{1,2\}}^{-1}(\mathbf{b}) \quad \forall c \in \mathbb{C}. \quad (2.8)$$

To generalise the vectorisation operation to higher-order tensors, we define for a tensor $\mathbf{A} \in \mathbb{C}^{N_1 \times \dots \times N_n}$ a *partial vectorisation* $\text{vec}_{\{i, \dots, j\}} : \mathbb{C}^{N_1 \times \dots \times N_n} \rightarrow \mathbb{C}^{N_1 N_i \dots N_j \times \dots \times N_n N_i \dots N_j}$ as the vectorisation of the subset of dimensions N_i, \dots, N_j of \mathbf{A} . As an example, the partial vectorisation of a tensor $\mathbf{C} \in \mathbb{C}^{N_1 \times N_2 \times N_3 \times N_4}$ with respect to the first and second dimensions is written as $\text{vec}_{\{1,2\}}(\mathbf{C}) \in \mathbb{C}^{N_3 N_1 N_2 \times N_4 N_1 N_2}$. Note that we here restrict ourselves to the special case where the vectorised dimensions are collapsed over all the remaining dimensions of the tensor rather than forming a separate, new dimension, which may not be the most general definition of partial vectorisation but sufficient for our purposes. For brevity, we write $\mathbf{A}_{\{i, \dots, j\}} \equiv \text{vec}_{\{i, \dots, j\}}(\mathbf{A})$. Note that we may equate the matrix vectorisation (2.4) with the partial vectorisation of a second-order tensor as $\text{vec}_{\{1,2\}}(\mathbf{A})$. The complete vectorisation of a tensor is then defined as the partial vectorisation of all dimensions, i.e.,

$$\text{vec}(\mathbf{A}) := \text{vec}_{\{1, \dots, n\}}(\mathbf{A}).$$

Similarly to (2.5), we may define a partial inverse vectorisation

$$\text{vec}_{\{i, \dots, j\}}^{-1}(\mathbf{A}_{\{i, \dots, j\}}) := \mathbf{A} \in \mathbb{C}^{N_1 \times \dots \times N_n}, \quad 1 \leq i \leq j \leq n. \quad (2.9)$$

2.1.3 The multidimensional discrete Fourier transform on matrix-vector form

In this section, we introduce a matrix-vector formulation of multilinear maps. The specific case of interest is the n -dimensional discrete Fourier transform (DFT), for which we introduce a version where a subset of the dimensions is transformed.

We first consider the case of a one-dimensional linear transform. Let $\mathcal{A} : \mathbb{C}^N \rightarrow \mathbb{C}^N$ be a complex-valued linear operator, and let $\boldsymbol{\psi} = [\psi_1 \ \cdots \ \psi_N]^\top \in \mathbb{C}^N$ be a signal vector. The \mathcal{A} -transform of $\boldsymbol{\psi}$ is then defined as $\mathcal{A}\{\boldsymbol{\psi}\}$. We may write the \mathcal{A} -transform in matrix form using an associated transformation matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$, such that $\mathcal{A}\{\boldsymbol{\psi}\} := \mathbf{A}\boldsymbol{\psi}$. In the n -dimensional case, the signal is represented by a tensor $\boldsymbol{\Psi} \in \mathbb{C}^{N_1 \times \cdots \times N_n}$, and the linear operator is defined as $\mathcal{A} : \mathbb{C}^{N_1 \times \cdots \times N_n} \rightarrow \mathbb{C}^{N_1 \times \cdots \times N_n}$. It was shown by [6] that the \mathcal{A} -transform of $\boldsymbol{\Psi}$ can be represented on matrix-vector form as the Kronecker product of lexicographically ordered 1D square transform matrices $\mathbf{A}_i \in \mathbb{C}^{N_i \times N_i}$, $i = 1, \dots, n$. Using the tensor vectorisation operators from the previous section, it is possible to write the transform as $\mathcal{A}\{\boldsymbol{\Psi}\} := \text{vec}^{-1}\{1, \dots, n\}((\mathbf{A}_1 \otimes \cdots \otimes \mathbf{A}_n)\text{vec}(\boldsymbol{\Psi}))$.

A particular case of the multilinear operator \mathcal{A} is the n -dimensional DFT $\mathcal{F}\{\boldsymbol{\Psi}\} := \text{vec}^{-1}_{\{1, \dots, n\}}((\mathbf{F}_1 \otimes \cdots \otimes \mathbf{F}_n)\text{vec}(\boldsymbol{\Psi}))$, where $\mathbf{F}_i \in \mathbb{C}^{N_i \times N_i}$ is the i th DFT matrix, for $i = 1, \dots, n$. The DFT matrices are Vandermonde matrices on the form

$$\mathbf{F}_i := \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & \omega_i & \cdots & \omega_i^{N_i-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_i^{N_i-1} & \cdots & \omega_i^{(N_i-1)^2} \end{bmatrix}, \quad (2.10)$$

with $\omega_i := \exp(-j2\pi/N_i)$ and inverse given by

$$\mathbf{F}_i^{-1} = \frac{1}{N_i} \mathbf{F}_i^H \quad (2.11)$$

Using the properties of the Kronecker product (2.3c)–(2.3d), the inverse DFT is also generalised to the n -dimensional case as

$$\begin{aligned} \mathcal{F}^{-1}\{\boldsymbol{\Psi}\} &= \text{vec}^{-1}_{\{1, \dots, n\}}\left((\mathbf{F}_1 \otimes \cdots \otimes \mathbf{F}_n)^{-1} \text{vec}(\boldsymbol{\Psi})\right) \\ &\stackrel{(2.3d)}{=} \text{vec}^{-1}_{\{1, \dots, n\}}\left(\left(\mathbf{F}_1^{-1} \otimes \cdots \otimes \mathbf{F}_n^{-1}\right) \text{vec}(\boldsymbol{\Psi})\right) \\ &\stackrel{(2.11)}{=} \text{vec}^{-1}_{\{1, \dots, n\}}\left(\left(\frac{1}{N_1} \mathbf{F}_1^H \otimes \cdots \otimes \frac{1}{N_n} \mathbf{F}_n^H\right) \text{vec}(\boldsymbol{\Psi})\right) \\ &\stackrel{(2.3e)}{=} \frac{1}{N_1 \cdots N_n} \text{vec}^{-1}_{\{1, \dots, n\}}\left(\left(\mathbf{F}_1^H \otimes \cdots \otimes \mathbf{F}_n^H\right) \text{vec}(\boldsymbol{\Psi})\right) \\ &\stackrel{(2.3c)}{=} \frac{1}{N_1 \cdots N_n} \text{vec}^{-1}_{\{1, \dots, n\}}\left(\left(\mathbf{F}_1 \otimes \cdots \otimes \mathbf{F}_n\right)^H \text{vec}(\boldsymbol{\Psi})\right) \\ &= \frac{1}{N_1 \cdots N_n} \mathcal{F}^H\{\boldsymbol{\Psi}\}, \end{aligned} \quad (2.12)$$

2. Background

where the H-operator denotes the Hermitian transpose. The n -dimensional DFT and its inverse transform all dimensions of the input tensor Ψ . If, however, we consider only a subset of k dimensions indexed by

$$\mathcal{D}_k = \{d_1, \dots, d_k : d_j \in \{1, \dots, n\} : d_j \neq d_l \forall j \neq l; k \leq n\}, \quad |\mathcal{D}_k| = k, \quad (2.13)$$

we may define a \mathcal{D}_k -dimensional DFT for the n -dimensional tensor Ψ as

$$\begin{aligned} \mathcal{F}_{\mathcal{D}_k} \{\Psi\} &:= \text{vec}_{\{1, \dots, n\}}^{-1} \left(\mathbf{F}_{\mathcal{D}_k} \text{vec}(\Psi) \right), \quad \text{where} \\ \mathbf{F}_{\mathcal{D}_k} &= \bigotimes_{j=1}^n \begin{cases} \mathbf{F}_{d_j} & \text{if } d_j \in \mathcal{D}_k, \\ \mathbf{I}_{N_{d_j}} & \text{otherwise} \end{cases}. \end{aligned} \quad (2.14)$$

Here, $\mathbf{I}_{N_{d_j}}$ denotes the $N_{d_j} \times N_{d_j}$ identity matrix. The inverse of the \mathcal{D}_k -dimensional DFT is similarly to (2.12) shown to be

$$\begin{aligned} \mathcal{F}_{\mathcal{D}_k}^{-1} \{\Psi\} &= \text{vec}_{\{1, \dots, n\}}^{-1} \left(\mathbf{F}_{\mathcal{D}_k}^{-1} \text{vec}(\Psi) \right) \\ &\stackrel{(2.3d)}{=} \text{vec}_{\{1, \dots, n\}}^{-1} \left(\left[\bigotimes_{j=1}^n \begin{cases} \mathbf{F}_{N_{d_j}}^{-1} & \text{if } d_j \in \mathcal{D}_k \\ \mathbf{I}_{N_{d_j}}^{-1} & \text{otherwise} \end{cases} \right] \text{vec}(\Psi) \right) \\ &\stackrel{(2.11)}{=} \text{vec}_{\{1, \dots, n\}}^{-1} \left(\left[\bigotimes_{j=1}^n \begin{cases} \frac{1}{N_{d_j}} \mathbf{F}_{d_j}^H & \text{if } d_j \in \mathcal{D}_k \\ \mathbf{I}_{N_{d_j}}^H & \text{otherwise} \end{cases} \right] \text{vec}(\Psi) \right) \\ &\stackrel{(2.3e)}{=} \frac{1}{N_{d_1} \cdots N_{d_k}} \text{vec}_{\{1, \dots, n\}}^{-1} \left(\left[\bigotimes_{j=1}^n \begin{cases} \mathbf{F}_{d_j}^H & \text{if } d_j \in \mathcal{D}_k \\ \mathbf{I}_{N_{d_j}}^H & \text{otherwise} \end{cases} \right] \text{vec}(\Psi) \right) \\ &\stackrel{(2.3c)}{=} \frac{1}{N_{d_1} \cdots N_{d_k}} \text{vec}_{\{1, \dots, n\}}^{-1} \left(\mathbf{F}_{\mathcal{D}_k}^H \text{vec}(\Psi) \right) \\ &= \frac{1}{N_{d_1} \cdots N_{d_k}} \mathcal{F}_{\mathcal{D}_k}^H \{\Psi\}. \end{aligned} \quad (2.15)$$

Both the n -dimensional DFT and the \mathcal{D}_k -dimensional DFT as defined with tensor notation are generalisations of the one-dimensional DFT, and may be normalised by the product of the dimensions of the indexed dimensions to become unitary transformations.

2.2 Massive MIMO and CSI acquisition

The development of multi-antenna systems in the mobile networks sector for close to three decades has successively led to more deployments of these technologies with each new generation of mobile networks, from a limited footprint in 3G systems to becoming an integral part of 4G, and subsequently having a global spread with massive MIMO solutions in the 5G mobile networks of today [7], [8]. There are three main multi-antenna techniques in massive MIMO, which together with the antenna arrays, including both radio hardware and software, make up the massive MIMO system: spatial multiplexing, beamforming,

and nullforming. With these techniques, network capacity, coverage and throughput can be significantly increased by multiplexing transmitted data streams over the same frequency-time resource with high directivity in the spatial domain, while simultaneously suppressing beam gain for other receivers than the intended UE [7].

The capabilities of massive MIMO are dependent upon the knowledge of the channel, the so-called CSI, at the BS in order to determine spatial filters for transmitted signals that align with the propagation channel, not only for DL but also for uplink (UL) reception. CSI acquisition for the receiver is achieved by letting the transmitter send known reference signals, also known as training pilots, which can be used for UE-side channel estimation and determining appropriate spatial filters for combining signals at the receiver [2, Ch. 6.4]. Channel estimation for downlink TX precoding and subsequently optimised downlink transmission is less tractable and requires either deriving the CSI at the receiver and sending this information back to the transmitter, referred to as closed-loop CSI feedback, or by performing channel measurements on the TX side and exploiting the reciprocity of the channel. The two approaches offers different capabilities and trade-offs, and are used interchangeably depending on the scenario and requirements [9]. In the closed-loop CSI feedback scheme, CSI-RSs are sent from the transmitter over different time-frequency resources, which allows the receiver to estimate the channel matrix between the TX and RX antenna pairs [2, Ch. 6.4]. There are many different ways to design the CSI feedback procedure, categorised as either explicit or implicit feedback. In explicit feedback, the full channel realisation or closely related properties are sent back to the transmitter. Implicit feedback, used in the 3GPP standards, comprises UE-side selection of TX transmission beam parameters from a finite set of configurations in a precoder codebook, based on properties of the estimated channel. For the other type of CSI acquisition based on UL sounding, the UE instead transmits reference signals, referred to as sounding reference signals (SRS), and the BS estimates the channel properties for TX precoding. This reciprocity-based beamforming works well for time division duplexing (TDD), where the frequency resource is shared between UL and DL transmission, and the channel then is reciprocal; in frequency division duplexing (FDD) however, there is no complete reciprocity and UL channel estimation is only suitable for long-term CSI [9]. Furthermore, for scenarios of poor UL coverage, closed-loop CSI feedback provides more stable performance compared to reciprocity-based CSI acquisition.

With the background for feedback-based CSI acquisition established, we can now present the system model and the associated channel model and antenna configuration that will form the basis for the application of the proposed methods in this thesis.

2.3 System model

For the purpose of this work, we consider a massive MIMO system at link level in a DL scenario, with BS as the transmitter and UE as the receiver. Both the BS and UE are equipped with uniform planar arrays (UPAs) with dual-polarised antenna element pairs, as illustrated in Figure 2.1. By 3GPP convention for CSI-RS port numbering and antenna port mapping, the spatial dimensions of UPAs are arranged as (N_v, N_h, N_p) , where N_v and N_h are the number of antenna elements along the vertical and horizontal spatial

2. Background




Figure 2.1: Antenna configuration of a massive MIMO system at link level with a BS (left) and a UE (right), both equipped with dual-polarised UPAs. The total number of TX antenna elements is $N_t = 4 \times 8 \times 2 = 64$ and the total number of RX antenna elements is $N_r = 1 \times 2 \times 2 = 4$.

dimensions, respectively, and N_p is the number of polarisations assumed to be identical for the transmitter and receiver [2, Ch. 8.3]. The total number of TX antenna elements is thus $N_t = N_{t,v} N_{t,h} N_p$, and the total number of RX antenna elements is $N_r = N_{r,v} N_{r,h} N_p$. In the orthogonal frequency-division multiplexing (OFDM) transmission scheme used in NR, a resource element consisting of one subcarrier over one OFDM symbol defines the fundamental waveform resource unit [8]. The subcarriers are spaced by $\Delta_f = 1/T$ such that the subcarriers are orthogonal to each other, where T together with a cyclic prefix duration becomes the OFDM symbol duration. For a given subcarrier k and at a specific instance in time, the relation between a transmitted signal $\mathbf{x}_k \in \mathbb{C}^{N_t}$ and the received signal $\mathbf{y}_k \in \mathbb{C}^{N_r}$ may be formulated as the linear system

$$\mathbf{y}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{e}_k, \quad (2.16)$$

where $\mathbf{H}_k \in \mathbb{C}^{N_r \times N_t}$ is the channel matrix, and $\mathbf{e}_k \in \mathbb{C}^{N_r}$ is an additive error term that contains both noise and interference [2, Ch. 5.3]. For the purpose of this work we assume ideal conditions, under which \mathbf{e} vanishes. In practice, the signal-to-interference-and-noise ratio (SINR) determines the quality in the estimate of the channel matrix \mathbf{H} from the received signal \mathbf{y} . The proposed methods in this thesis are however not limited to ideal conditions, but can be applied to any scenario where the DL channel can be estimated.

The two main approaches to represent the channel matrix \mathbf{H}_k are as a *directional propagation channel* and a *radio channel* [2, Ch. 3.6]. In the directional propagation channel model, the channel is described in terms of the radio wave propagation paths between the transmitter and receiver, whereas the radio channel model directly relates the transmitted and received signals. Due to the large number of antenna elements in massive MIMO systems and the need for spatial multiplexing to achieve high data rates, the directional propagation channel model is often preferred as it provides a more detailed description of the channel. We here therefore introduce a directional propagation model for the DL channel for OFDM-based transmission in MIMO systems [2]. The properties that characterise the i th propagation path are defined to be $\{\tau_i, \theta_{t,i}, \varphi_{t,i}, \theta_{r,i}, \varphi_{r,i}, a_i, \boldsymbol{\Psi}_i\}$. The propagation delay τ_i is the time it takes for the signal to travel from the BS to the UE along the i th path, and is related to the path distance d_i by $\tau_i = d_i/c$. The transmitted waves depart from the BS in the direction $(\theta_{t,i}, \varphi_{t,i})$ defined by the azimuth and zenith

2. Background

angles $\varphi_{t,i}$ and $\theta_{t,i}$, respectively. The UE receives the waves with a complex path gain a_i in analogously the direction $(\theta_{r,i}, \varphi_{r,i})$. Alternatively, the angle of departure (AoD) can be described by the unit vector $\widehat{\mathbf{r}}_{t,i}$, and the angle of arrival (AoA) by the unit vector $\widehat{\mathbf{r}}_{r,i}$. The wave vectors are then given by $\mathbf{k}_{t,i} = (2\pi f_c/c)\widehat{\mathbf{r}}_{t,i}$ and $\mathbf{k}_{r,i} = (2\pi f_c/c)\widehat{\mathbf{r}}_{r,i}$, where f_c is the carrier frequency, c is the speed of light and $\widehat{\mathbf{r}}_{t,i}, \widehat{\mathbf{r}}_{r,i}$ denote the unit normal radial TX and RX. Note that $\mathbf{k}_{r,i}$ is defined in the opposite direction of the propagating wave. With Jones representation form, the transmitted wave along the i th with an electric field strength $\mathbf{E}_{t,i} = [E_{\theta,t,i} \ E_{\varphi,t,i}]^\top$ is received at the UE with an electric field strength of $\mathbf{E}_{r,i} = a_i e^{-j2\pi f_c \tau_i} \boldsymbol{\Psi}_i \mathbf{E}_{t,i}$, where $e^{-j2\pi f_c \tau_i}$ is the phase shift, f_c is the carrier frequency, and $\boldsymbol{\Psi}_i$ is the polarisation scattering matrix given by

$$\boldsymbol{\Psi}_i := \begin{bmatrix} \psi_{\theta_t \theta_r, i} & \psi_{\varphi_t \theta_r, i} \\ \psi_{\theta_t \varphi_r, i} & \psi_{\varphi_t \varphi_r, i} \end{bmatrix}. \quad (2.17)$$

Here, $\psi_{\theta_t \theta_r, i}$ represents the relative change in amplitude and phase shift of the $\widehat{\boldsymbol{\theta}}_t$ -polarised wave along the $\widehat{\boldsymbol{\theta}}_r$ -direction at the receiver, and similarly for the other coefficients. The polarisation scattering matrix $\boldsymbol{\Psi}_i$ is normalised such that the diagonal elements are equal to one.

With multiple TX and RX antenna element-pairs, a number of intricacies follow in order to accurately model the channel for the dual-polarised UPAs. The polarisations for the p th TX antenna element and the q th RX antenna element are expressed as

$$\widehat{\boldsymbol{\psi}}_{t,p}(\theta, \varphi) = \begin{cases} \widehat{\boldsymbol{\psi}}_{t,A}(\theta, \varphi) & \text{if } p = 1, \dots, N_t/2, \\ \widehat{\boldsymbol{\psi}}_{t,B}(\theta, \varphi) & \text{if } p = N_t/2 + 1, \dots, N_t, \end{cases} \quad (2.18)$$

and

$$\widehat{\boldsymbol{\psi}}_{r,q}(\theta, \varphi) = \begin{cases} \widehat{\boldsymbol{\psi}}_{r,A}(\theta, \varphi) & \text{if } q = 1, \dots, N_r/2, \\ \widehat{\boldsymbol{\psi}}_{r,B}(\theta, \varphi) & \text{if } q = N_r/2 + 1, \dots, N_r, \end{cases} \quad (2.19)$$

where the polarisation vectors for polarisation A are given by $\widehat{\boldsymbol{\psi}}_{r,A}(\theta, \varphi) = [\psi_{r,A,\theta} \ \psi_{r,A,\varphi}]^\top$ and $\widehat{\boldsymbol{\psi}}_{t,A}(\theta, \varphi) = [\psi_{t,A,\theta} \ \psi_{t,A,\varphi}]^\top$, respectively, and analogously for polarisation B . In effect, four polarisation coefficients are defined for each antenna element pair, from which we create a polarisation coefficient matrix for the i th propagation path

$$\boldsymbol{\chi}_i := \begin{bmatrix} \widehat{\boldsymbol{\psi}}_{r,A}^\top(\theta_{r,i}, \varphi_{r,i}) \boldsymbol{\Psi}_i \widehat{\boldsymbol{\psi}}_{t,A}(\theta_{t,i}, \varphi_{t,i}) & \widehat{\boldsymbol{\psi}}_{r,A}^\top(\theta_{r,i}, \varphi_{r,i}) \boldsymbol{\Psi}_i \widehat{\boldsymbol{\psi}}_{t,B}(\theta_{t,i}, \varphi_{t,i}) \\ \widehat{\boldsymbol{\psi}}_{r,B}^\top(\theta_{r,i}, \varphi_{r,i}) \boldsymbol{\Psi}_i \widehat{\boldsymbol{\psi}}_{t,A}(\theta_{t,i}, \varphi_{t,i}) & \widehat{\boldsymbol{\psi}}_{r,B}^\top(\theta_{r,i}, \varphi_{r,i}) \boldsymbol{\Psi}_i \widehat{\boldsymbol{\psi}}_{t,B}(\theta_{t,i}, \varphi_{t,i}) \end{bmatrix}. \quad (2.20)$$

The positions of the TX and RX antenna elements are modelled by the vectors $\{\mathbf{d}_{t,p} : p = 1, \dots, N_t\}$ and $\{\mathbf{d}_{r,q} : q = 1, \dots, N_r\}$, respectively, such that the vectors are relative to the local origin(s) of the propagation paths. Given that the receiver travels at a velocity \mathbf{v}_r , the position of the q th RX antenna element at time t is assumed to be $\mathbf{d}_{r,q} + \mathbf{v}_r t$. For simplification, let the complex amplitude patterns $g_t(\theta, \varphi)$ and $g_r(\theta, \varphi)$ be constant across antenna arrays for the TX and RX antennas, respectively. With this setup, we construct array response vectors for the TX and RX antennas of dimensions $N_t/2 \times 1$ and $N_r/2 \times 1$, respectively, as

2. Background

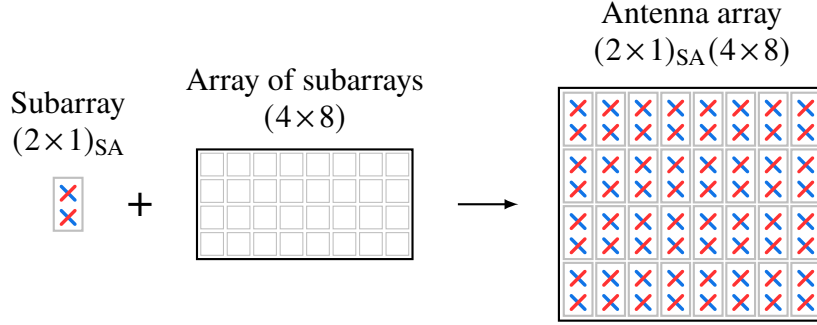


Figure 2.2: Illustration of a $(4 \times 1)_{SA} (2 \times 8)$ UPA, divided into SAs with one RF chain per polarisation and SA.

$$\mathbf{a}_{t/2}(\theta_{t,i}, \varphi_{t,i}) = g_t(\theta_{t,i}, \varphi_{t,i}) \left[e^{j\mathbf{d}_{t,1}^T \mathbf{k}_{t,i}} \quad \dots \quad e^{j\mathbf{d}_{t,N_t/2}^T \mathbf{k}_{t,i}} \right]^T, \quad (2.21)$$

$$\mathbf{a}_{r/2}(\theta_{r,i}, \varphi_{r,i}) = g_r(\theta_{r,i}, \varphi_{r,i}) \left[e^{j\mathbf{d}_{r,1}^T \mathbf{k}_{r,i}} \quad \dots \quad e^{j\mathbf{d}_{r,N_r/2}^T \mathbf{k}_{r,i}} \right]^T. \quad (2.22)$$

The MIMO channel for the k th subcarrier at the time instance t may then be expressed as

$$\mathbf{H}_k = \sum_i \lambda_{i,k}(t) \chi_i \otimes \left(\mathbf{a}_{r/2}(\theta_{r,i}, \varphi_{r,i}) \mathbf{a}_{t/2}^T(\theta_{t,i}, \varphi_{t,i}) \right), \quad (2.23)$$

where $\lambda_{i,k} := \alpha_i e^{j2\pi f_{D,i} t} e^{-j2\pi k \Delta_f \tau_i}$ is introduced for compact notation, with $\alpha_i \in \mathbb{C}$ being the attenuation and phase shifts along the i th propagation path, and $f_{D,i}$ the Doppler frequency such that $2\pi f_{D,i} = \mathbf{v}_r^T \mathbf{k}_{r,i}$ [2, Eq. 5.51]. The Kronecker product results in a block structure of the channel matrix;

$$\mathbf{H}_k = \begin{bmatrix} \mathbf{H}_{k,1,1} & \mathbf{H}_{k,1,2} \\ \mathbf{H}_{k,2,1} & \mathbf{H}_{k,2,2} \end{bmatrix}, \quad (2.24)$$

where each block

$$\mathbf{H}_{k,l,m} = \sum_i \lambda_{i,k}(t) [\chi_i]_{l,m} \mathbf{a}_{r/2}(\theta_{r,i}, \varphi_{r,i}) \mathbf{a}_{t/2}^T(\theta_{t,i}, \varphi_{t,i}) \quad (2.25)$$

corresponds to a MIMO channel with single-polarised antenna elements at the transmitter and receiver.

UEs are often assumed to be positioned in a limited range of AoDs, making it possible to design beamforming vectors at the TX side that are either fixed or vary at a slow rate for down-tilt steering. This is done by grouping the antenna elements into subarrays (SAs) as depicted in Figure 2.2, each consisting of a few antenna elements that are connected to a single radio frequency (RF) chain (per polarisation). The resulting antenna virtualisation forms an array of subarrays (AoSA) that appears as a dual-polarised UPA, but where the number of virtual antenna ports \tilde{N}_t are equal to the number of RF chains. Using the notation of [2], we denote an AoSA of $\tilde{N}_v \times \tilde{N}_h$ SAs, where each SA is of size $M_v \times M_h$, as $(M_v \times M_h)_{SA} (\tilde{N}_v \times \tilde{N}_h)$. The choice of M_v and M_h is an important parameter for an effective antenna system and depends on the application and placement of BSs. We formalise the SA beamforming for the k th subcarrier as a weight matrix $\mathbf{W}_{SA,k}$ of size $N_t \times \tilde{N}_t$, implemented at RF in close proximity to the antenna elements, resulting in an

2. Background

effective DL channel $\mathbf{H}_{\text{AoSA},k} := \mathbf{H}_k \mathbf{W}_{\text{SA},k} \in \mathbb{C}^{N_r \times \tilde{N}_t}$. We will henceforth only consider virtual UPAs at the transmitter, and consequently only the effective DL channel matrix; for simplicity, we let N_t denote the number of virtual antenna ports.

In multilayer MIMO transmission, multiple signal layers or symbol streams are spatially multiplexed and beamformed towards the same or different UEs, known as single-user MIMO (SU-MIMO) and multi-user MIMO (MU-MIMO), respectively. In this procedure, a digital precoding matrix $\mathbf{W}_k \in \mathbb{C}^{N_t \times L}$ is applied to the transmitted symbols $\mathbf{s}_k \in \mathbb{C}^L$ at the baseband, where L is the number of spatial layers or symbol streams, such that the transmitted signal is given by $\mathbf{x}_k = \mathbf{W}_k \mathbf{s}_k$. It is necessary, but not sufficient, that the number of layers $L \leq \min\{N_t, N_r\}$. The precoding is applied to subbands in the frequency domain with a granularity of 12 consecutive subcarriers, known as physical resource blocks (PRBs) [8]. Because of this design, the DL channel \mathbf{H}_k is estimated at the UE over the (one-dimensional) PRBs. A set of N_T consecutive estimates of the DL channel by the UE thus forms a tensor $\hat{\mathbf{H}}$ of size $N_r \times N_t \times N_{\text{PRB}} \times N_T$. Note however that, due to current system limitations of OFDM transmission in NR, a single channel estimate $\hat{\mathbf{H}}$ may be constructed in from OFDM symbols across time. Since we disregard time-varying properties of the channel in this work, this aspect is not further discussed. For an overview of the different 3GPP precoding schemes for massive MIMO, see [10]; an in-depth discussion of the various precoding techniques can be found in [2].

In the next chapter, we will first develop general methods for near-optimal sparse sampling of source signals before applying these to the problem of feedback-based CSI acquisition in massive MIMO systems.

2. Background

3

Sparse sampling of CSI-RS using submodular optimisation

This chapter addresses the sparse pilot-placement problem by establishing a methodology for constructing near-optimal sparse sounding patterns for CSI-RS that maximises the recoverability of the full-rank DL channel. We first cast the problem as a general linear inverse problem in which the linear map is to be estimated from measurements when the source signal is sparsely sampled; this formulation is then generalised for tensor-valued linear systems, allowing for the sparse sampling across multiple dimensions of the signal tensor. The search for optimal sampling schemes is then formalised as an optimisation problem, which for objective functions of certain properties allows for greedy algorithms with near-optimality guarantees. To quantify the quality of any candidate sampling pattern, we briefly introduce frame theory and the p -frame potential that quantify the source correlation in physical systems, and subsequently propose a cost function based on this scalar function. This lets us deduce an improved optimality of the greedy algorithm, and then extend these methods to real-world scenarios with stochastic systems. Finally, we briefly apply our findings to sparse-sampling in massive MIMO feedback CSI, showing how near-optimal CSI-RS placement across antenna ports, frequency resources and/or time slots can be structured.

3.1 Sparse sampling for linear inverse problems

In the following work, we will consider the noiseless linear inverse problem

$$\mathbf{y} = \Psi \mathbf{x}, \quad (3.1)$$

where $\mathbf{y} \in \mathbb{C}^N$ is the observed signal at the receiver, $\Psi \in \mathbb{C}^{N \times M}$ is the forward linear map representing the physical system, and $\mathbf{x} \in \mathbb{C}^M$ is the transmitted signal. Contrary to the standard inverse problem, we assume that the source signal \mathbf{x} is known and at our disposal, and that we are tasked with estimating the physical system Ψ . More specifically, we here consider the case where the signal \mathbf{x} may be transmitted in a L -sparse manner, i.e., only a subset of $L \leq M$ of the M entries of \mathbf{x} are nonzero:

Definition 3.1 (Sampling sparsity.)

A signal $\mathbf{x} \in \mathbb{C}^M$ is said to be L -sparse if it has L nonzero entries, i.e., if the ℓ_0 -pseudonorm $\|\mathbf{x}\|_0 := \sum_{i=1}^M |x_i|^0 = L$, with the definition $0^0 \equiv 0$.

3. Sparse sampling of CSI-RS using submodular optimisation

$$\mathcal{L} = \{ 1, \quad 3, 4, \quad 7, \quad 9, 10, 11 \}$$

$$\varphi_{\mathcal{L}}^{\top} = \begin{array}{cccccccccccc} \blacksquare & \blacksquare & \square & \square & \square & \square & \square & \square & \square & \square & \square & \square \end{array}$$

Figure 3.1: A sampling scheme with a sampling set \mathcal{L} and the corresponding sampling vector $\varphi_{\mathcal{L}}$. Black and white cells represent zero and identity, respectively.

Sparsity enables data compression by storing only the indices and values of nonzero elements, avoiding redundancy. When data lacks sparsity, a *sparse representation* in some basis may often be identified, for example the angular (Fourier) domain, that enables compression in the original domain – we will revisit this topic in Chapter 4. Given the set of indices of the signal \mathbf{x} as $\mathcal{N} := \{1, \dots, M\}$, we may characterise the sparse sampling of \mathbf{x} by the set of indices $\mathcal{L} := \{i_1, \dots, i_L\} \subseteq \mathcal{N}$, with $|\mathcal{L}| = L$, that corresponds to the selected entries of \mathbf{x} to be transmitted. Equivalently, the sampling may be thought of as a *muting* of the source signal elements, $\mathcal{S} = \mathcal{L} \setminus \mathcal{N}$, as will become apparent in subsequent chapters; in this context, we will refer to *muting level* as the relative magnitude of the source signal muting: $|\mathcal{S}|/|\mathcal{N}| = 1 - |\mathcal{L}|/|\mathcal{N}|$. We define the sampling vector as an L -hot binary vector $\varphi_{\mathcal{L}} \in \{0, 1\}^M \setminus \{0\}^M$ such that the i th element of $\varphi_{\mathcal{L}}$ is given by

$$\varphi_{\mathcal{L},i} = \begin{cases} 1 & \text{if } i \in \mathcal{L}, \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in \mathcal{N}. \quad (3.2)$$

An illustrative example of a sampling scheme with a sampling set and the corresponding sampling vector is shown in Figure 3.1. Using this representation of sparsely sampling for the source vector \mathbf{x} , we introduce the sampling matrix as a diagonal matrix $\Phi_{\mathcal{L}} := \text{diag}(\varphi_{\mathcal{L}}) \in \{0, 1\}^{M \times M}$. Furthermore, we define the sparse physical system as $\Psi_{\mathcal{L}} := \Psi \Phi_{\mathcal{L}} \in \mathbb{C}^{N \times M}$ and let $\Psi'_{\mathcal{L}} \in \mathbb{C}^{N \times L}$ be the pruned matrix obtained by selecting the sampled (nonzero) columns of Ψ . The sparsely sampled linear inverse problem (3.1) can then be expressed as

$$\mathbf{y}_{\mathcal{L}} = \Psi \Phi_{\mathcal{L}} \mathbf{x} = \Psi_{\mathcal{L}} \mathbf{x} = \Psi'_{\mathcal{L}} \mathbf{x}_{\mathcal{L}}, \quad (3.3)$$

where $\mathbf{y}_{\mathcal{L}}$ is the measured signal for the sampling scheme and $\mathbf{x}_{\mathcal{L}}$ is the pruned signal vector obtained by selecting the sampled entries of \mathbf{x} . Figure 3.2 illustrates this reformulation for the sampling scheme found in Figure 3.1.

The linear inverse problem (3.1) may be generalised to the case of tensors. Let $\mathbf{Y} \in \mathbb{C}^{N \times S_1 \times \dots \times S_s}$ be the sensor signal tensor, $\mathbf{X} \in \mathbb{C}^{M \times S_1 \times \dots \times S_s}$ the source signal tensor, and $\Psi \in \mathbb{C}^{N \times M \times S_1 \times \dots \times S_s}$ the forward linear map representing the physical system, with shared dimensions S_1, \dots, S_s for the source and sensor signals (e.g., time and/or frequency). Then, the linear inverse problem may be expressed as the matrix equation

$$Y_{ik_1 \dots k_s} = \Psi_{ijk_1 \dots k_s} X_{jk_1 \dots k_s}, \quad (3.4)$$

where $Y_{ik_1 \dots k_s}$ is the $(ik_1 \dots k_s)$ th entry of the sensor signal tensor \mathbf{Y} , and similarly for \mathbf{X} and Ψ . Alternatively, we can employ tensor vectorisation as defined in Section 2.1.2 to formulate the linear inverse problem for the complete tensors as

$$\text{vec}(\mathbf{Y}) = \Psi_{(1, \dots, s)} \text{vec}(\mathbf{X}), \quad (3.5)$$

3. Sparse sampling of CSI-RS using submodular optimisation




Figure 3.2: A sampling scheme for the sparsely sampled linear inverse problem (3.3) with a sampling set \mathcal{L} and corresponding sampling matrix $\Phi_{\mathcal{L}}$. Colored cells indicate values of corresponding entries of the quantity.

with $\text{vec}(\mathbf{Y})$ and $\text{vec}(\mathbf{X})$ vectors of dimensions $NS_1 \cdots S_s$ and $MS_1 \cdots S_s$, respectively. The partially vectorised tensor $\Psi_{(1, \dots, s)} \equiv \text{vec}_{(1, \dots, s)}(\Psi)$ of dimensions $NS_1 \cdots S_s \times MS_1 \cdots S_s$ is the matrix obtained by vectorising the shared dimensions $S_1 \cdots S_s$ of (3.4). With this formulation, the generalised problem (3.5) is equivalent to the linear inverse problem (3.1) with the tensor $\Psi_{(1, \dots, s)}$ as the forward linear map, and consequently, a sampling scheme for (3.5) can be defined with a sampling set $\mathcal{L} \subseteq \mathcal{N}$ with $|\mathcal{N}| = MS_1 \cdots S_s$ and $|\mathcal{L}| = L$. The partial vectorisation of the physical system Ψ effectively contracts all dimensions over which the source signal \mathbf{x} is sampled into a single dimension, which will be necessary for solving the multi-dimensional sampling problem.

3.2 Submodular optimisation for designing near-optimal sampling schemes

The aim that we set for the sparse sampling problem (3.3) is to estimate the full-rank physical system Ψ from the sampled measurements $\mathbf{Y}_{\mathcal{L}}$. This is a non-trivial task, as the sampling scheme \mathcal{L} makes the problem underdetermined. Finding an optimal solution to this ill-posed problem requires two steps; choosing a sampling scheme \mathcal{L} and a reconstruction method for the physical system Ψ . Two questions immediately arise from this observation:

1. What is the optimal sampling scheme \mathcal{L} to use for the problem that minimises the error in the estimate of Ψ ?
2. What are good metrics to use for the error in the estimate of Ψ , or equivalently the quality of reconstruction of the physical system?

The answer to the first question requires solving the combinatorial optimisation problem

$$(O3.1a) \quad \begin{aligned} & \arg \max_{\mathcal{L} \subset \mathcal{N}} f(\mathcal{L}) \\ & \text{subject to } |\mathcal{L}| = L, \end{aligned}$$

where the objective function f is a metric that quantifies accuracy whereby the full physical system Ψ can be reconstructed from the sampled measurements $\mathbf{Y}_{\mathcal{L}}$. This is a constrained cardinality problem [11] for the known number of selected samples L , which is nonconvex in \mathcal{L} . With the related sampling vector, the problem can equivalently be expressed as the

3. Sparse sampling of CSI-RS using submodular optimisation

boolean optimisation problem

$$(O3.1b) \quad \begin{aligned} & \arg \max_{\varphi_{\mathcal{L}}: \mathcal{L} \subset \mathcal{N}} f(\varphi_{\mathcal{L}}) \\ & \text{subject to } \|\varphi_{\mathcal{L}}\|_1 = L, \end{aligned}$$

with the accuracy metric f here defined as a function of the sampling vector $\varphi_{\mathcal{L}}$ and $\|\cdot\|_1$ denoting the ℓ_1 -norm. For our purposes, we will henceforth refer to the set formulation of the optimisation problem (O3.1a). The exhaustive search for the optimal sampling scheme \mathcal{L}^* to (O3.1a)–(O3.1b) is computationally infeasible for large M , since it entails evaluating the objective function f for all $\binom{M}{L}$ possible combinations of L indices from the set \mathcal{N} . Branch-and-bound methods yield an exact solution to the problem [12], [13], but are also prohibitively expensive to compute and unrealistic for large M . The practical approach for solving (O3.1a)–(O3.1b) consists of using approximate algorithms that yield a suboptimal solution to the problem. One such method is convex relaxation of the boolean problem to the continuous domain $\{0, 1\}^M \rightarrow [0, 1]^M$ and using thresholding or randomised rounding to obtain a solution in the discrete domain [14], [15]. Another approximation technique that also answers the second question, and which we will use in this work, is to find a function f for the accuracy of the reconstruction of the physical system Ψ that is *submodular*, i.e., a function that satisfies the property of diminishing returns. Formally, submodularity is defined as follows:

Definition 3.2 (Submodular function.)

A set function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ defined for all subsets of \mathcal{N} is submodular if $\forall \mathcal{X}, \mathcal{Y} \subseteq \mathcal{N}$

$$f(\mathcal{X} \cup \mathcal{Y}) + f(\mathcal{X} \cap \mathcal{Y}) \leq f(\mathcal{X}) + f(\mathcal{Y}). \quad (3.4)$$

A more intuitive definition was provided by Nemhauser *et al.* [16] that captures the notion of diminishing returns:

Definition 3.3 (Submodular function [16].)

A set function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is submodular if $\forall \mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{N}$ and $i \in \mathcal{N} \setminus \mathcal{Y}$

$$f(\mathcal{X} \cup \{i\}) - f(\mathcal{X}) \geq f(\mathcal{Y} \cup \{i\}) - f(\mathcal{Y}). \quad (3.5)$$

Notably, submodularity is closed under nonnegative linear combinations [17], i.e., if f_1 and f_2 are submodular functions, then $f = \alpha f_1 + \beta f_2$ is also submodular for $\alpha, \beta \geq 0$. We define two more properties for set functions:

Definition 3.4 (Monotone function.)

A set function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is monotone if $\forall \mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{N}$

$$f(\mathcal{X}) \leq f(\mathcal{Y}). \quad (3.6)$$

Definition 3.5 (Normalised set function.)

A set function $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ defined for all subsets of \mathcal{N} is normalised if $f(\emptyset) = 0$.

Nemhauser *et al.* [16] proved that a submodular function that is also monotone and normalised can be greedily maximised with near-optimality guarantees. In such a greedy algorithm, shown in Algorithm 3.1, one iteratively chooses the element that maximises

3. Sparse sampling of CSI-RS using submodular optimisation

Algorithm 3.1 Greedy algorithm for submodular function maximisation.

Input: \mathcal{N} set of elements, $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ submodular, monotone and normalised, L number of elements to select.

$\mathcal{L} = \emptyset$

for $l = 1, \dots, L$ **do**

$$\hat{i} = \min \left(\arg \max_{i \in \mathcal{N} \setminus \mathcal{L}} f(\mathcal{L} \cup \{i\}) \right)$$

$\mathcal{L} = \mathcal{L} \cup \{\hat{i}\}$

end for

return \mathcal{L}

the marginal gain of the objective function. For uniqueness of the algorithm, the smallest integer element is selected in cases where several indices yield the maximal function value. In addition, no other algorithm that runs in polynomial time is guaranteed to find a solution that is better than $(1 - 1/e)$ of the optimal solution in polynomial, unless $P = NP$ [18]. We here present the near-optimality of Algorithm 3.1 for monotone, normalised submodular functions as formulated by Ranieri *et al.* [19]:

Theorem 3.1 (Near-optimal maximalisation of submodular function [16].)

Let $f: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a monotone, normalised submodular set function defined over all subsets of a finite set \mathcal{N} . Let $\hat{\mathcal{L}} \subset \mathcal{N}$ be the set of L elements chosen by Algorithm 3.1, and let

$$\mathcal{L}^* := \text{lex min} \left(\arg \max_{\mathcal{A} \subset \mathcal{N}: |\mathcal{A}|=L} f(\mathcal{A}) \right) \quad (3.7)$$

be the lexicographically smallest optimal set of L elements. Then

$$f(\hat{\mathcal{L}}) \geq \left(1 - \frac{1}{e}\right) f(\mathcal{L}^*), \quad (3.8)$$

where e is Euler's number.

We will now proceed to extend Algorithm 3.1 to suit stochastic systems and realizations of the physical system in datasets.

3.3 The generalised average-case criterion for stochastic systems

Thus far, we have only discussed sampling schemes for the linear inverse problem (3.1) in the context of fixed linear operators Ψ . In practice however, the physical system Ψ is to be considered a stochastic variable that is not known *a priori*. The full physical system Ψ can be observed, or estimated to some extent, from a finite number of measurements \mathbf{y} or through modelling of the physical system. The result is a set of K realisations $\{\Psi^{(k)}\}_{k=1}^K$ of the physical system over which we want to design an optimal sparse sampling scheme.

3. Sparse sampling of CSI-RS using submodular optimisation

The problem setup can then be formulated as a generalised average-case criterion [20] of the sampling scheme over the K realisations of the physical system;

$$(O3.2) \quad \begin{aligned} & \arg \max_{\mathcal{L} \subset \mathcal{N}} \sum_{k=1}^K w_k f^{(k)}(\mathcal{L}) \\ & \text{subject to } |\mathcal{L}| = L, \end{aligned}$$

where $f^{(k)}$ is the objective function for the k th realisation of the physical system Ψ and $w_k \geq 0$ for $k=1, \dots, K$, are weights that can be used to adjust the relative contribution of each realisation. An important result is that the weighted average of a set of functions preserves normality, monotonicity and submodularity, as shown in the following lemma.

Lemma 3.1 (Closedness of normality, monotonicity and submodularity under weighted average.)

Let $f^{(k)} : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a set function for $k=1, \dots, K$. If $f^{(k)}$ is normalised, monotone and submodular for all k , then for weights $\mathbf{w} = \{w_k\}_{k=1}^K$ such that $w_k \geq 0$ for all k , the weighted average

$$\mathbb{E}_{\mathbf{w}}[f(\mathcal{L})] = \sum_{k=1}^K w_k f^{(k)}(\mathcal{L}), \quad (3.9)$$

is also normalised, monotone and submodular.

Proof. The proof follows directly from the definitions of normality, monotonicity and submodularity. Trivially, the weighted average is normalised such that $\mathbb{E}_{\mathbf{w}}[f(\emptyset)] = 0$. For monotonicity and submodularity, we note that the weighted average is a linear combination of the set functions $f^{(k)}(\mathcal{L})$ for each realisation k in the dataset with factors w_k . Since monotonicity and submodularity are closed under nonnegative linear combinations, we conclude that $\mathbb{E}_{\mathbf{w}}[f^{(k)}(\mathcal{L})]$ is also monotone and submodular. \square

It follows directly from Theorem 3.1 and Lemma 3.1 that a near-optimal greedy algorithm is also valid for the generalised average-case criterion (O3.2). In particular, we denote the unweighted average of the set functions, with $w_k = 1/K$ for $k=1, \dots, K$, as

$$\mathbb{E}[f(\mathcal{L})] := \frac{1}{K} \sum_{k=1}^K f^{(k)}(\mathcal{L}). \quad (3.10)$$

3.4 The p -frame potential

To characterise sampling schemes for the linear inverse problem (3.1), and subsequently find near-optimal sampling schemes, we will introduce ideas from the theory of frames and the work of Ranieiri *et al.* [19]. This subject was originally conceptualised in the context of non-harmonic Fourier series by Duffin and Schaeffer [21] in the early 1950s, and has become significant in signal processing as well as in non-uniform sampling theory [22], [23]. A frame is formally defined as a collection of finite-dimensional elements $\{\mathbf{x}_i\}_{i=1}^M$ in a Hilbert space \mathcal{H} as follows [22]:

3. Sparse sampling of CSI-RS using submodular optimisation

Definition 3.6 (Frames.)

Let \mathcal{H} be a separable Hilbert space with inner product $\langle \mathbf{x}, \mathbf{y} \rangle$ and norm $\|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{x} \rangle^{1/2}$. A sequence $\{\mathbf{x}_i\}_{i=1}^M$ is a frame for \mathcal{H} if there exist frame bounds $A, B > 0$ such that

$$\forall \mathbf{y} \in \mathcal{H}, \quad A\|\mathbf{y}\|^2 \leq \sum_{n=1}^M |\langle \mathbf{y}, \mathbf{x}_i \rangle|^2 \leq B\|\mathbf{y}\|^2. \quad (3.11)$$

A frame is said to be tight if $A = B$.

Frames are not necessarily orthonormal bases (ONBs) in \mathcal{H} , but can nonetheless be more advantageous than ONBs for noise reduction [22, Ch. 1.2]. The Hilbert space in this work will henceforth be the complex vector space \mathbb{C}^N , and the finite frame under consideration consists of the column vectors $\{\boldsymbol{\psi}_i\}_{i=1}^M$ of the linear map representing the physical system $\boldsymbol{\Psi} \in \mathbb{C}^{N \times M}$ that forms the so-called frame matrix. The frame property (3.11) may then alternatively be expressed in terms of the physical system matrix [24]:

$$\forall \mathbf{y} \in \mathbb{C}^N, \quad A\|\mathbf{y}\|_2^2 \leq \|\boldsymbol{\Psi}^H \mathbf{y}\|_2^2 \leq B\|\mathbf{y}\|_2^2, \quad A, B > 0. \quad (3.12)$$

To characterise a sequence of finite frames using a scalar potential energy function, Benedetto and Fickus [23] introduced the concept of frame potential (FP):

Definition 3.7 (Frame potential.)

The frame potential (FP) of a frame matrix $\boldsymbol{\Psi} \in \mathbb{C}^{N \times M}$ is defined as

$$\text{FP}(\boldsymbol{\Psi}) := \sum_{i=1}^M \sum_{j=1}^M |\langle \boldsymbol{\psi}_i, \boldsymbol{\psi}_j \rangle|^2, \quad (3.13)$$

where $\langle \cdot, \cdot \rangle$ is the inner product in \mathbb{C}^N .

This functional encourages orthogonality between the frame vectors $\boldsymbol{\psi}_i$ and $\boldsymbol{\psi}_j$ for $i \neq j$; the more orthogonal the frame vectors are, the lower the FP [19]. The FP is minimised for tight frames and may thus be thought of as a measure of tightness of $\boldsymbol{\Psi}$ [25]. Note that the definition we have provided here differs from the definition of Ranieiri *et al.* [19], since we are interested in sampling of sources rather than sensors as in the usual case and we therefore effectively treat the frame potential of the transposed system $(\boldsymbol{\Psi} \boldsymbol{\Phi}_L)^\top = \boldsymbol{\Phi}_L \boldsymbol{\Psi}^\top$ as compared to [19]. This however does not yield any loss of generality with respect to frame theory, since we treat $\{\boldsymbol{\psi}_i\}_{i=1}^M$ as a finite frame.

The special case of finite unit-norm tight frames (FUNTFs) where the frame elements $\boldsymbol{\psi}_i$ have unit norm is of particular interest. These frames are placed on the finite-dimensional unit sphere and exhibit Parseval-like relations that resemble orthonormal bases [23]. We here however do not restrict ourselves to FUNTFs, but rather consider the more general case of finite frames without unit-norm vectors for which the FP still is well-defined [23].

Ehler and Okoudjou [26] extended the FP functional to a p -frame potential for $0 < p < \infty$, valid also for complex frames:

3. Sparse sampling of CSI-RS using submodular optimisation

Definition 3.8 (*p*-frame potential.)

The *p*-frame potential (*p*-FP) of a system $\Psi \in \mathbb{C}^{N \times M}$ for $0 < p < \infty$ is defined as

$$\text{FP}_p(\Psi) := \sum_{i=1}^M \sum_{j=1}^M |\langle \psi_i, \psi_j \rangle|^p \quad 0 < p < \infty, \quad (3.14)$$

where ψ_i is the *i*th column of the matrix Ψ and $\langle \cdot, \cdot \rangle$ is the inner product in \mathbb{C}^N .

An interesting property of the *p*-FP is that it can be expressed in terms of the Gramian $\mathbf{G} = \Psi^H \Psi$ [19], [27] as

$$\text{FP}_p(\Psi) = \sum_{i=1}^M \sum_{j=1}^M |G_{ij}|^p, \quad (3.15)$$

where G_{ij} is the (*i*, *j*)th element of \mathbf{G} . In particular for $p=2$, the ordinary FP can be written as

$$\text{FP}_2(\Psi) = \sum_{i=1}^M \sum_{j=1}^M |G_{ij}|^2 = \|\mathbf{G}\|_F^2 = \text{tr}(\mathbf{G}\mathbf{G}^H) = \sum_{i=1}^M |\lambda_i|^2, \quad (3.16)$$

where $\|\cdot\|_F$ is the Frobenius norm and λ_i are the eigenvalues of the Gramian. In communication theory, the average of the Gramian is referred to as the TX covariance matrix [9]. If the physical system is sparsely sampled with a scheme \mathcal{L} , the Gramian of that system is

$$\mathbf{G}_{\mathcal{L}} = \Psi_{\mathcal{L}}^H \Psi_{\mathcal{L}} = \Psi^H \Phi_{\mathcal{L}} \Psi \quad (3.17)$$

and its average is the TX covariance matrix with precoding matrix $\Phi_{\mathcal{L}}$. As a consequence, the FP is known as the total summed correlation in communication networks [28]. It is possible to extend the notion of frame potential to multi-dimensional systems using our formulation in Section 3.1, where the Gramian is computed across the sampled dimensions.

3.5 Related work

With the above introduction to greedy algorithms for near-optimal sparse sampling schemes and to the FP, a review of related work is in place to articulate the contributions of this thesis work in relation to the literature. Ranieiri *et al.* [19] introduced a submodular cost function to be optimised based on the FP for sensor placement design and a number of efforts have been made since to further develop the proposed algorithm. This submodular objective function was modified to a weighted version of the FP by [29]. It was shown by [30] that the originally proposed greedy backward elimination algorithm can be replaced by forward elimination suitable for submodular optimisation. A framework for sparse sampling of tensors was developed by [31] and extended the submodular optimisation algorithm and corresponding FP-based objective function to the case of multidomain graph signals, also applying their method to sparse MIMO receivers. Further applications of the method of Ranieiri *et al.* include wireless sensor networks [32] and MIMO radar imaging [33], [34]. In the field of massive MIMO, submodular optimization for antenna selection has been studied with the channel capacity [35]–[37] as a basis for the objective

3. Sparse sampling of CSI-RS using submodular optimisation

function, but to the best of the author's knowledge the FP has not been applied in this context, and in particular to the problem of overhead reduction for CSI-RS.

In the following, we make a number of new contributions based on the original work of Ranieiri *et al.* [19]: A submodular optimisation algorithm is proposed with an objective function that is normalised between zero and one and based on the p -FP in complex vector space, improving the near-optimality guarantee of Nemhauser *et al.* [16] with respect to the objective function and to the p -FP. By proving normality, monotonicity and submodularity of this new cost function for all $p \in (0, \infty)$, setting $p = 1$ yields a computationally more efficient procedure than earlier versions of the algorithm. For practical application in the case of stochastic systems, the algorithm is also formulated for the case of learning from a dataset by applying the average-case criterion from the previous section. Furthermore, this setup extends the sparse sampling optimisation approach to placement of *source* signals. Finally, using the framework for sampling of multilinear systems in Section 3.1, a generalisation to higher dimensions is achieved.

3.6 A submodular function based on the p -frame potential

We will now present an objective function for the problem of designing optimal sampling schemes for source signals (O3.1a) using Algorithm 3.1. The p -FP is not submodular with respect to sparse sampling schemes; instead, we similar to [19] define an objective function using the complement of the selected source signal elements $\mathcal{S} := \mathcal{N} \setminus \mathcal{L}$:

$$G_p(\mathcal{S}) := 1 - \frac{\text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{S}})}{\text{FP}_p(\Psi)}. \quad (3.18)$$

This proposed cost function computes the ratio between the p -FP of the physical system with sampled elements $\mathcal{L} = \mathcal{N} \setminus \mathcal{S}$, and the p -FP of the full-rank physical system matrix. We formally show that G_p is normalised, monotone and submodular in Theorem 3.2.

Theorem 3.2 (Normality, monotonicity and submodularity of the objective function.)
The function $G_p(\mathcal{S})$ as defined in (3.18) is normalised, monotone, and submodular for $0 < p < \infty$.

Proof. The following proof extends the proof given by [19]. First, $G_p(\mathcal{S})$ is trivially shown to be normalised such that $G_p(\emptyset) = 0$ and $G_p(\mathcal{N}) = 1$, since $\Psi_{\mathcal{N}} \equiv \Psi$ and the p -frame potential of the zero-matrix $\text{FP}_p(\Psi_{\emptyset}) = \text{FP}_p(\mathbf{0}) = 0$. To prove that G_p is monotone, we assume that $\mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{N}$ and $i \in \mathcal{N} \setminus \mathcal{Y}$. Then, by using a change of variables (*) $\mathcal{A} := \mathcal{N} \setminus \mathcal{X}$,

$$\begin{aligned} G_p(\mathcal{X} \cup \{i\}) - G_p(\mathcal{X}) &= \frac{\text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{X}})}{\text{FP}_p(\Psi)} - \frac{\text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{X} \cup \{i\}})}{\text{FP}_p(\Psi)} \\ &= c(\text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{X}}) - \text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{X} \cup \{i\}})) \\ &\stackrel{(*)}{=} c \left(\sum_{n, m \in \mathcal{A} \cup \{i\}} |\langle \psi_n, \psi_m \rangle|^p - \sum_{n, m \in \mathcal{A}} |\langle \psi_n, \psi_m \rangle|^p \right) \end{aligned}$$

3. Sparse sampling of CSI-RS using submodular optimisation

$$= 2c \left(|\langle \boldsymbol{\psi}_i, \boldsymbol{\psi}_i \rangle|^p + \sum_{n \in \mathcal{A}} |\langle \boldsymbol{\psi}_n, \boldsymbol{\psi}_i \rangle|^p \right) \geq 2c \cdot 0 = 0, \quad (3.19)$$

where $c = 1/\text{FP}_p(\boldsymbol{\Psi})$ is a positive real constant. To show that G_p is submodular, we assume without loss of generality that $\mathcal{Y} = \mathcal{X} \cup \{j\}$ and get

$$\begin{aligned} & G_p(\mathcal{X} \cup \{i\}) - G_p(\mathcal{X}) - [G_p(\mathcal{Y} \cup \{i\}) - G_p(\mathcal{Y})] \\ &= G_p(\mathcal{X} \cup \{i\}) - G_p(\mathcal{X}) - [G_p(\mathcal{X} \cup \{i, j\}) - G_p(\mathcal{X} \cup \{j\})] \\ &= \frac{1}{\text{FP}_p(\boldsymbol{\Psi})} [\text{FP}_p(\boldsymbol{\Psi}_{\mathcal{A} \cup \{i, j\}}) - \text{FP}_p(\boldsymbol{\Psi}_{\mathcal{A} \cup \{i\}}) - \text{FP}_p(\boldsymbol{\Psi}_{\mathcal{A} \cup \{j\}}) + \text{FP}_p(\boldsymbol{\Psi}_{\mathcal{A}})] \\ &= c \left[2 \sum_{n \in \mathcal{A} \cup \{j\}} |\langle \boldsymbol{\psi}_n, \boldsymbol{\psi}_i \rangle|^p - 2 \sum_{n \in \mathcal{A}} |\langle \boldsymbol{\psi}_n, \boldsymbol{\psi}_i \rangle|^p \right] \\ &= 2c |\langle \boldsymbol{\psi}_j, \boldsymbol{\psi}_i \rangle|^p \geq 0. \end{aligned} \quad (3.20)$$

□

The main consequence of this result is that the degree p of the FP may be chosen arbitrarily. We may thus choose $p = 1$ to our advantage of minimising the computational complexity of the problem and the numerical stability when computing the frame potential. In the rest of the work, we will denote $G_1 \equiv G$ for the case $p = 1$. A greedy algorithm for the sampling scheme (O3.1a) can be constructed by using the cost function G_p as the objective function is shown in Algorithm 3.2, which iteratively removes elements from the set \mathcal{N} that maximises the marginal gain of the cost function G_p until the desired number of elements are left. Secondly, since $G_p \in [0, 1]$, we can further improve the optimality estimate of the greedy algorithm in Theorem 3.1:

Corollary 3.1 (Near-optimality of the greedy algorithm for the normalised cost function.)

Let $G_p : 2^{\mathcal{N}} \rightarrow [0, 1]$ be a normalised, monotone and submodular set function defined over all subsets of a finite set \mathcal{N} and for all $0 < p < \infty$. Let $\hat{\mathcal{S}} \subset \mathcal{N}$ be the complement to the set of L elements chosen by Algorithm 3.2, and let

$$\mathcal{S}^* := \text{lex min} \left(\arg \max_{\mathcal{A} \subset \mathcal{N}: |\mathcal{A}| = |\mathcal{N}| - L} G_p(\mathcal{A}) \right) \quad (3.21)$$

be the complement to the lexicographically smallest optimal set of L elements with respect to G_p . Then

$$G_p(\hat{\mathcal{S}}) \geq \max \left\{ G_p(\hat{\mathcal{S}}), 1 - \frac{1}{e} \right\} G_p(\mathcal{S}^*). \quad (3.22)$$

Proof. It follows directly from the definition of G_p (3.18) and Theorem 3.1 that $G_p(\hat{\mathcal{S}}) \geq (1 - 1/e)G_p(\mathcal{S}^*)$. Because $G_p \in [0, 1]$, we also have that $G_p(\hat{\mathcal{S}})G_p(\mathcal{S}^*) \leq G_p(\hat{\mathcal{S}})$.

3. Sparse sampling of CSI-RS using submodular optimisation

Algorithm 3.2 Greedy algorithm for the sampling scheme.

Input: \mathcal{N} set of elements, $G_p : 2^{\mathcal{N}} \rightarrow [0, 1]$ according to (3.18), L number of elements to select.
 $\mathcal{S} = \emptyset$
for $l = 1, \dots, |\mathcal{N}| - L$ **do**
 $\hat{i} = \min \left(\arg \max_{i \in \mathcal{N} \setminus \mathcal{S}} G_p(\mathcal{S} \cup \{i\}) \right)$
 $\mathcal{S} = \mathcal{S} \cup \{\hat{i}\}$
end for
 $\mathcal{L} = \mathcal{N} \setminus \mathcal{S}$
return \mathcal{L}

Therefore, the lower bound on $G_p(\hat{\mathcal{S}})$ becomes

$$G_p(\hat{\mathcal{S}}) \geq \max \left\{ G_p(\hat{\mathcal{S}})G_p(\mathcal{S}^*), \left(1 - \frac{1}{e}\right)G_p(\mathcal{S}^*) \right\}, \quad (3.23)$$

which proves the result. \square

Corollary 3.1 provides a near-optimal guarantee of the greedy algorithm with respect to the cost function G_p , but does not imply any result in terms of the frame potential. We will therefore show that Algorithm 3.2 also gives an upper bound to the p -FP in the following corollary:

Corollary 3.2 (Near-optimality of the greedy algorithm with respect to the p -FP.)

Let $\hat{\mathcal{L}} \subset \mathcal{N}$ be the set of L elements chosen by Algorithm 3.2, and let

$$\mathcal{L}^* := \text{lex min} \left(\arg \max_{\mathcal{A} \subset \mathcal{N}; |\mathcal{A}| = |\mathcal{N}| - L} G_p(\mathcal{A}) \right) \quad (3.24)$$

be the lexicographically smallest optimal set of L elements. Then, for the p -frame potential $\text{FP}_p : \mathbb{C}^{N \times M} \rightarrow [0, \infty)$,

$$\text{FP}_p(\Psi_{\hat{\mathcal{L}}}) \leq \gamma \text{FP}_p(\Psi) + (1 - \gamma) \text{FP}_p(\Psi_{\mathcal{L}^*}) \quad (3.25)$$

with $\gamma = \min \left\{ \frac{1}{e}, \frac{\text{FP}_p(\Psi_{\hat{\mathcal{L}}})}{\text{FP}_p(\Psi)} \right\}$.

Proof. Using the definition of G_p (3.18), the result of Corollary 3.1 directly yields

$$1 - \frac{\text{FP}_p(\Psi_{\mathcal{N} \setminus \hat{\mathcal{S}}})}{\text{FP}_p(\Psi)} \geq (1 - \gamma) \left(1 - \frac{\text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{S}^*})}{\text{FP}_p(\Psi)} \right); \quad \gamma = \min \left\{ \frac{1}{e}, \frac{\text{FP}_p(\Psi_{\mathcal{N} \setminus \hat{\mathcal{S}}})}{\text{FP}_p(\Psi)} \right\}, \quad (3.26)$$

where $\hat{\mathcal{S}}$ is the complement of the L elements selected by Algorithm 3.2 and \mathcal{S}^* the complement of the optimal set of L elements with respect to G_p (see Corollary 3.1).

3. Sparse sampling of CSI-RS using submodular optimisation

Algorithm 3.3 Greedy algorithm for the generalised average-case criterion sampling scheme.

Input: \mathcal{N} set of elements, K realisations of the physical system Ψ , $G_p^{(k)}: 2^{\mathcal{N}} \rightarrow [0, 1]$ according to (3.18) for $k = 1, \dots, K$, $\mathbf{w} = \{w_k\}_{k=1}^K$ weights for the K realisations, and L number of elements .

$\mathcal{S} = \emptyset$

for $l = 1, \dots, |\mathcal{N}| - L$ **do**

$$\hat{i} = \min \left(\arg \max_{i \in \mathcal{N} \setminus \mathcal{S}} \mathbb{E}_{\mathbf{w}} \left[G_p^{(k)}(\mathcal{S} \cup \{i\}) \right] \right)$$

$\mathcal{S} = \mathcal{S} \cup \{\hat{i}\}$

end for

$\mathcal{L} = \mathcal{N} \setminus \mathcal{S}$

return \mathcal{L}

Rewriting the expression gives

$$\text{FP}_p(\Psi_{\mathcal{N} \setminus \hat{\mathcal{S}}}) \leq \gamma \text{FP}_p(\Psi) + (1 - \gamma) \text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{S}^*}). \quad (3.27)$$

Under the change of variables $\mathcal{L} = \mathcal{N} \setminus \mathcal{S}$, it is clear that $\text{FP}_p(\Psi_{\mathcal{N} \setminus \hat{\mathcal{S}}}) = \text{FP}_p(\Psi_{\hat{\mathcal{L}}})$, and that the optimisation problem

$$(O3.3a) \quad \arg \min_{\mathcal{L} \subset \mathcal{N}} \text{FP}_p(\mathcal{L}) \quad \text{subject to} \quad |\mathcal{L}| = L$$

is equivalent to

$$(O3.3b) \quad \arg \min_{\mathcal{S} \subset \mathcal{N}} \text{FP}_p(\mathcal{S}) \quad \text{subject to} \quad |\mathcal{S}| = |\mathcal{N}| - L$$

for all $0 < p < \infty$. From this follows that $\text{FP}_p(\Psi_{\mathcal{N} \setminus \mathcal{S}^*}) = \text{FP}_p(\Psi_{\mathcal{L}^*})$, which concludes the proof. \square

We may also extend these results for the generalised average-case criterion (O3.2), since $\mathbb{E}_{\mathbf{w}}[G_p(\mathcal{L})]$ is also normalised, monotone and submodular according to Lemma 3.1. We can therefore apply the greedy algorithm Algorithm 3.1 to the generalised average-case criterion with analogous optimality guarantees as in Corollaries 3.1 and 3.2, explicitly shown in Algorithm 3.3.

3.7 Sparse sampling of the DL channel

With the greedy algorithms for sparse sampling of linear multi-dimensional inverse problems developed in the previous sections, we now formulate the necessary steps to apply the method to the problem of sparse CSI-RS sounding in a massive MIMO DL channel in accordance with the procedures described in [3], [38]. Let $\mathbf{H} \in \mathbb{C}^{N_r \times N_t \times N_{\text{PRB}} \times N_{\text{T}}}$ be a time series of a massive MIMO downlink channel estimated by the UE. Since the

3. Sparse sampling of CSI-RS using submodular optimisation

downlink channel is a tensor, we will in the following use partial tensor vectorisation as described in Section 3.1 to transform linear system of downlink transmission to two dimensions

$$\text{vec}(\mathbf{Y}) = \mathbf{H}_{(\text{PRB}, \text{T})} \text{vec}(\mathbf{X}), \quad (3.26)$$

where $\mathbf{Y} \in \mathbb{C}^{N_t \times N_{\text{PRB}} \times N_T}$ is the received signal at the UE, $\mathbf{X} \in \mathbb{C}^{N_t \times N_{\text{PRB}} \times N_T}$ is the transmitted CSI pilot signal from the BS, and $\mathbf{H}_{(\text{PRB}, N_T)} \in \mathbb{C}^{N_r \times N_{\text{PRB}} \times N_T \times N_t \times N_{\text{PRB}} \times N_T}$ is the partially vectorised channel matrix over the PRBs and the time instances. A sparsely sounded CSI-RS across antenna ports in a sampling scheme \mathcal{L} then results in an effective channel matrix

$$\mathbf{H}_{\mathcal{L}, (\text{PRB}, \text{T})} = \mathbf{H}_{(\text{PRB}, \text{T})} (\boldsymbol{\Phi}_{\mathcal{L}} \otimes \mathbf{I}_{N_{\text{PRB}}} \otimes \mathbf{I}_{N_T}), \quad (3.27)$$

where $\boldsymbol{\Phi}_{\mathcal{L}}$ is the $N_t \times N_t$ sampling matrix over the transmitter antenna ports. The sparse sampling scheme may however not only be confined to the antenna ports, but also include the subbands;

$$\mathbf{H}_{\mathcal{L}, (\text{PRB}, \text{T})} = \mathbf{H}_{(\text{PRB}, \text{T})} \left(\left(\boldsymbol{\Phi}_{\mathcal{L}_1} \oplus \cdots \oplus \boldsymbol{\Phi}_{\mathcal{L}_{N_{\text{PRB}}}} \right) \otimes \mathbf{I}_T \right), \quad \mathcal{L} = \bigcup_{i=1}^{N_{\text{PRB}}} \mathcal{L}_i, \quad (3.28)$$

with sampling matrices $\boldsymbol{\Phi}_{\mathcal{L}_i} \in \{0, 1\}^{N_t \times N_t}$ for $i = 1, \dots, N_{\text{PRB}}$. If the downlink channel has been measured over a number of time instances N_T , one may also consider sounding CSI-RS differently over time, i.e., the sampling scheme may yield a sparse downlink channel matrix

$$\mathbf{H}_{\mathcal{L}, (\text{PRB}, \text{T})} = \mathbf{H}_{(\text{PRB}, \text{T})} \left(\bigoplus_{i=1}^{N_{\text{PRB}}} \bigoplus_{j=1}^{N_T} \boldsymbol{\Phi}_{\mathcal{L}_{i,j}} \right), \quad \mathcal{L} = \bigcup_{i=1}^{N_{\text{PRB}}} \bigcup_{j=1}^{N_T} \mathcal{L}_{i,j}, \quad (3.29)$$

where each sampling matrix $\boldsymbol{\Phi}_{\mathcal{L}_{i,j}}$ has dimensions $N_t N_{\text{PRB}} N_T \times N_t N_{\text{PRB}} N_T$.

The algorithms presented in Section 3.6 should be thought of as naïve if applied to the downlink channel estimation problem for all possible sampling schemes $\mathcal{L} \subset \mathcal{N}$, since they are purely based on computing the frame potential of all CSI-RS sounding elements and do not integrate any domain knowledge or constraints on the sampling scheme. A more sophisticated approach, although not explored in this work, is to limit the set of all elements \mathcal{N} to subsets that are feasible for the requirements of CSI-RS, e.g., to sparsely sample only a subset of antenna ports, PRBs, and/or time instances.

The question still remains for which measures are appropriate for the information contained in, or equivalently, the capability for reconstruction of the full-rank channel in sparse sampling scheme. In the next chapter, we will therefore cover the *decompression* of sparsely sampled systems, with particular focus on neural network-based methods for reconstruction of the full channel from sparsely sounded CSI-RS.

3. Sparse sampling of CSI-RS using submodular optimisation

4

Neural network-based reconstruction of sparsely sounded CSI-RS

Sparse sampling schemes for linear inverse problems such as the massive MIMO downlink channel estimation problem allows for compression of the full system from a limited set of transmissions. On the receiver side, the challenge lies in accurately recovering the full channel state information (CSI) from these compressed measurements. This chapter will therefore cover methods for reconstruction of sparsely sampled linear inverse problems, with focus on a PINN for reconstruction of the massive MIMO downlink channel from sparsely sounded CSI-RS. Before the ANN architecture is presented, we will formulate the reconstruction problem for sparsely sampled linear inverse problems, briefly review some classical algorithm from CS theory, and propose a divergence measure based on the beamspace channel power spectrum.

4.1 Reconstruction of sparsely sampled linear inverse problems

Given a linear inverse problem on the form (3.1), we assume that the source signal has been sparsely sampled with a sampling scheme \mathcal{L} such that the measurements $\mathbf{y}_{\mathcal{L}} \in \mathbb{C}^N$ yield a column-sparse system $\mathbf{\Psi}_{\mathcal{L}} = \mathbf{\Psi}\Phi_{\mathcal{L}} \in \mathbb{C}^{N \times M}$ of the full physical system $\mathbf{\Psi} \in \mathbb{C}^{N \times M}$. Reconstructing $\mathbf{\Psi}$ from the sparse measurements by a learned operator $f_{\theta} : \mathbb{C}^{N \times M} \rightarrow \mathbb{C}^{N \times M}$ with a set of parameters θ can be formulated as the recovery problem

$$(O4.1) \quad \begin{aligned} & \arg \min_{\theta} \quad \left\| \widehat{\mathbf{\Psi}} - \mathbf{\Psi} \right\|_F^2 \\ & \text{subject to} \quad \widehat{\mathbf{\Psi}} = f_{\theta}(\mathbf{\Psi}_{\mathcal{L}}). \end{aligned}$$

When $\mathbf{\Psi}$ is not *a priori* known, we assume instead K realisations $\{\mathbf{\Psi}^{(k)}\}_{k=1}^K$ of the physical system and replace the objective function with the mean squared error (MSE) or the normalised mean squared error (NMSE) over the K measurements, here defined to be

$$\text{MSE}(\widehat{\mathbf{\Psi}}, \mathbf{\Psi}) := \mathbb{E} \left[\left\| \widehat{\mathbf{\Psi}} - \mathbf{\Psi} \right\|_F^2 \right]; \quad \text{NMSE}(\widehat{\mathbf{\Psi}}, \mathbf{\Psi}) := \mathbb{E} \left[\frac{\left\| \widehat{\mathbf{\Psi}} - \mathbf{\Psi} \right\|_F^2}{\left\| \mathbf{\Psi} \right\|_F^2} \right]. \quad (4.1)$$

The recovery problem (O4.1) and the associated reconstruction measures are trivially generalised to the case of multi-dimensional linear systems described in Section 3.1.

4. Neural network-based reconstruction of sparsely sounded CSI-RS

4.2 Classical compressed sensing-based reconstruction algorithms

The reconstruction problem (O4.1) has been studied extensively in the literature and is a well-known problem in the field of CS, which both addresses the problem of sampling design and sparse recovery algorithms. Originally developed by Candès *et al.* [39] and Donoho [40], CS theory aims to reconstruct sparse signals sampled below the Nyquist rate and comprises both compression (sampling) and acquisition (reconstruction) of the signal. The theory is based on the assumption that the signal to be reconstructed is sparse in some basis [39], [40], i.e., it can be represented as a linear combination of a small number of basis functions, and hence is compressible. Relevant examples are signals in communication systems that often are sparse in the angular domain and thus can be compressed in a Fourier basis. For brevity, we will consider only a few of the most common CS algorithms and refer to [41] for a more thorough overview of CS-based reconstruction algorithms.

Let $\boldsymbol{\psi} \in \mathbb{C}^N$ be the vector to be reconstructed from a L -sparse set of measurements $\boldsymbol{\psi}_{\mathcal{L}} \in \mathbb{C}^N$ for a sampling scheme \mathcal{L} . We assume there exists a linear transformation \mathcal{A} with associated transformation matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$ that transforms the vector $\boldsymbol{\psi}$ into a sparse representation $\boldsymbol{\theta} := \mathcal{A}\{\boldsymbol{\psi}\} = \mathbf{A}\boldsymbol{\psi} \in \mathbb{C}^N$. If the linear transform is invertible, we can reconstruct the vector $\boldsymbol{\psi}$ from the sparse representation $\boldsymbol{\theta}$ as $\boldsymbol{\psi} = \mathcal{A}^{-1}\{\boldsymbol{\theta}\} = \mathbf{A}^{-1}\boldsymbol{\theta}$. The sparsely sampled vector may then be written on the form

$$\boldsymbol{\psi}_{\mathcal{L}} = \boldsymbol{\Phi}_{\mathcal{L}}\boldsymbol{\psi} = \boldsymbol{\Phi}_{\mathcal{L}}\mathcal{A}^{-1}\{\boldsymbol{\theta}\}. \quad (4.2)$$

We can now formulate the sparse reconstruction problem as an optimisation problem

$$(O4.2) \quad \begin{aligned} & \arg \min_{\boldsymbol{\theta}} \quad \|\boldsymbol{\theta}\|_0 \\ & \text{subject to} \quad \boldsymbol{\Phi}_{\mathcal{L}}\mathcal{A}^{-1}\{\boldsymbol{\theta}\} = \boldsymbol{\psi}_{\mathcal{L}}, \end{aligned}$$

where the goal is to find a feasible so-called minimum weight solution to the linear equations. This combinatorial optimisation problem is an example of a minimum relevant variables in linear system (Min-RVLS) problem [42, MP5] and difficult to solve because of the nonconvexity of the ℓ_0 -pseudonorm. In practice therefore, the sparse recovery problem (O4.2) is either solved with greedy algorithms or relaxed to a convex optimisation problem. Examples of greedy algorithms used include matching pursuit (MP) [43], the improved orthogonal matching pursuit (OMP) [44], and variations of these algorithms. Another approach is convex relaxation of the combinatorial problem [45], for which the ℓ_1 -relaxation has been shown to be the best substitute for $\|\cdot\|_0$; for example the result holds for $\|\cdot\|_p^p$ for all p [46] or can be explained geometrically [47]. This leads to the following modification:

$$(O4.3) \quad \begin{aligned} & \arg \min_{\boldsymbol{\theta}} \quad \|\boldsymbol{\theta}\|_1 \\ & \text{subject to} \quad \boldsymbol{\Phi}_{\mathcal{L}}\mathcal{A}^{-1}\{\boldsymbol{\theta}\} = \boldsymbol{\psi}_{\mathcal{L}} \end{aligned}$$

This convex optimisation problem is known as basis pursuit (BP), originally presented

4. Neural network-based reconstruction of sparsely sounded CSI-RS

by [48], and can be formulated as a linear program (LP) on the standard form

$$(O4.4) \quad \begin{aligned} & \min_z \quad \mathbf{1}^\top \mathbf{z} \\ & \text{subject to} \quad \Lambda \mathbf{z} = \boldsymbol{\psi}_{\mathcal{L}}, \\ & \quad \quad \quad \Lambda = [\boldsymbol{\Phi}_{\mathcal{L}} \mathbf{A}^{-1} \quad -\boldsymbol{\Phi}_{\mathcal{L}} \mathbf{A}^{-1}], \\ & \quad \quad \quad \mathbf{z} \geq 0, \end{aligned}$$

where \geq denotes the elementwise greater or equal to operator and $\mathbf{1}_N \in \{1\}^N$ is a vector of ones of length N . The solution to (O4.4) is given by the vector $\mathbf{z}^* = [\mathbf{u}^* \quad \mathbf{v}^*]$ and $\boldsymbol{\theta}^* = \mathbf{u}^* - \mathbf{v}^*$ solves (O4.3) [40]. In the case where Gaussian noise is present in the measurements, a relaxation of the equality constraint in the BP problem was proposed by [48] – the exact decomposition of $\boldsymbol{\psi}$ is in this case not desirable but rather a noiseless approximation of the signal. This modified BP problem is known as basis pursuit denoising (BPDN) [49] and is here given on the form

$$(O4.5) \quad \begin{aligned} & \arg \min_{\boldsymbol{\theta}} \quad \|\boldsymbol{\theta}\|_1 \\ & \text{subject to} \quad \|\boldsymbol{\Phi}_{\mathcal{L}} \mathcal{A}^{-1}\{\boldsymbol{\theta}\} - \boldsymbol{\psi}_{\mathcal{L}}\|_2 < \varepsilon, \end{aligned}$$

for some $\varepsilon > 0$. The BPDN problem is an ℓ_1 -relaxation of the empirical atomic decomposition method of [49], [50]. Because the residual norm $\|\boldsymbol{\Phi}_{\mathcal{L}} \mathcal{A}^{-1}\{\boldsymbol{\theta}\} - \boldsymbol{\psi}_{\mathcal{L}}\|_2$ does not account for non-sampled antenna ports, the standard BPDN problem (O4.5) may yield a solution where which for the corresponding non-sampled antenna ports are zero. To prevent this, penalty weights may be introduced or, as suggested by [46] in the context of MIMO channel estimation, a weighted ℓ_1 -norm can be used instead in the BPDN problem. In this work however, we will only consider the standard BPDN problem (O4.5).

To extend (O4.5) to the case of a tensor for the physical system $\boldsymbol{\Psi} \in \mathbb{C}^{N \times M \times S_1 \times \dots \times S_s}$ as introduced in Section 3.1, we define its sparse representation as $\boldsymbol{\theta} := \mathcal{A}\{\boldsymbol{\Psi}\}$ for a linear transformation \mathcal{A} . Using the concatenated sampling vector across the non-sampled dimensions

$$\tilde{\boldsymbol{\varphi}}_{\mathcal{L}} := \underbrace{[\boldsymbol{\varphi}_{\mathcal{L}}^\top \quad \dots \quad \boldsymbol{\varphi}_{\mathcal{L}}^\top]^\top}_{N \text{ times}} \in \{0, 1\}^{NMS_1 \dots S_s}, \quad (4.3)$$

we formulate the BPDN problem using the Hadamard (elementwise) product \odot as

$$(O4.6) \quad \begin{aligned} & \arg \min_{\boldsymbol{\theta}} \quad \|\text{vec}(\boldsymbol{\theta})\|_1 \\ & \text{subject to} \quad \left\| \left(\tilde{\boldsymbol{\varphi}}_{\mathcal{L}} \odot \text{vec}(\mathcal{A}^{-1}\{\boldsymbol{\theta}\}) \right) - \text{vec}(\boldsymbol{\Psi}_{\mathcal{L}}) \right\|_2 < \varepsilon. \end{aligned}$$

4. Neural network-based reconstruction of sparsely sounded CSI-RS

4.3 Compressed sensing-based reconstruction of sparsely sampled DL channel

From the theoretical background of sparse recovery algorithms in the previous section, we will now apply the BPDN problem (O4.5) to the downlink channel estimation problem. A time series of the DL channel $\mathbf{H} \in \mathbb{C}^{N_r \times N_t \times N_{\text{PRB}} \times T}$ is assumed to be sparse in the angular domain of the TX antenna ports. The transformation of the channel into TX beamspace is achieved by a discrete Fourier transform (DFT) over the spatial placement of the TX ports, which for a UPA with $N_{t,v}$ and $N_{t,h}$ antenna elements in the two spatial dimensions is given by $\mathbf{F}_{N_t} = \mathbf{F}_{N_{t,v}} \otimes \mathbf{F}_{N_{t,h}} \otimes \mathbf{I}_{N_p} \in \mathbb{C}^{N_t \times N_t}$, where $\mathbf{F}_{N_{t,v}}$ and $\mathbf{F}_{N_{t,h}}$ are the ordinary DFT matrices for the two spatial dimensions of the TX ports, and \mathbf{I}_{N_p} is the identity matrix for the polarisation dimension. This enables us to write the Fourier-transformed DL channel \mathbf{H} as follows:

$$\mathbf{H}_{\text{bs}} = \mathcal{F}_{\{N_t\}}\{\mathbf{H}\} = \text{vec}^{-1}\left(\left(\mathbf{I}_{N_r} \otimes \mathbf{F}_{N_t} \otimes \mathbf{I}_{N_{\text{PRB}}} \otimes \mathbf{I}_T\right) \text{vec}(\mathbf{H})\right). \quad (4.4)$$

The BPDN problem (O4.5) can then be formulated for the sparsely sampled DL channel $\mathbf{H}_{\mathcal{L}}$ from Section 3.7 accordingly as

$$(O4.7) \quad \begin{aligned} & \arg \min_{\boldsymbol{\theta}} \|\text{vec}(\boldsymbol{\theta})\|_1 \\ & \text{subject to} \quad \left\| \left(\tilde{\boldsymbol{\varphi}}_{\mathcal{L}} \odot \text{vec}\left(\mathcal{F}_{\{N_t\}}^{-1}\{\boldsymbol{\theta}\}\right) \right) - \text{vec}(\mathbf{H}_{\mathcal{L}}) \right\|_2 < \varepsilon, \end{aligned}$$

with $\tilde{\boldsymbol{\varphi}}_{\mathcal{L}}$ as defined in (4.3).

4.4 A distance measure for the reconstruction problem with sparsity in the angular domain

To quantify the reconstruction performance of the sparse downlink channel estimation problem in a meaningful manner, we have to critically examine the hitherto presented metrics in Section 4.3. A good performance metric in this sense, to answer the question posed in Section 4.1, should be based on the sparse beamspace representation of the downlink channel where information is concentrated in a small number of discrete AoDs. The MSE and NMSE (4.1) metrics are in this light not ideal candidates, since it is clear from their definitions that they are invariant to unitary transformation and in particular to the DFT with unitary normalisation. A more suitable approach is to instead analyse the power spectrum, or power spectral density (PSD), of the downlink channel. For our purposes, we will define the PSD here as the one-dimensional probability mass function $p_{\mathbf{H}}(i) := p_i / \|\mathbf{p}\|_1$ with $\mathbf{p} = |\text{vec}(\mathbf{H}_{\text{bs}})|^{\circ 2}$ and \circ denoting the elementwise (Hadamard) power. The spectral density of the complete downlink channel to be reconstructed is considered *a priori* information. There are several divergence measures commonly used for comparing power spectral densities with origin in probability theory, for example the Kullback-Leibler divergence [51], [52] and the Hellinger distance [53], [54]. Itakura and Saito [55], [56] introduced another distance measure that has become one of the most frequently used in

4. Neural network-based reconstruction of sparsely sounded CSI-RS

speech signal processing [57]. We here use the definition of Enqvist and Karlsson [54] of the Itakura-Saito distance, with the slight modification of introducing an additive term $\varepsilon \ll 1$ for the purpose of making the definition well-defined for our purposes:

$$D_{\text{IS}}(p_{\widehat{\mathbf{H}}}, p_{\mathbf{H}}) := \left\langle \frac{p_{\widehat{\mathbf{H}}} + \varepsilon}{p_{\mathbf{H}} + \varepsilon}, \mathbf{1} \right\rangle - \left\langle \log \frac{p_{\widehat{\mathbf{H}}} + \varepsilon}{p_{\mathbf{H}} + \varepsilon}, \mathbf{1} \right\rangle - |R_{\mathbf{H}}|. \quad (4.5)$$

Here, $p_{\widehat{\mathbf{H}}}$ and $p_{\mathbf{H}}$ are the power spectrums of the reconstructed sparsely sampled downlink channel $\widehat{\mathbf{H}}$ with sampling scheme \mathcal{L} and the full downlink channel \mathbf{H} , respectively; the term $|R_{\mathbf{H}}|$ denotes the length of the support R of the PSD for \mathbf{H} . The measure is equal to zero only if the two distributions are identical, $p_{\widehat{\mathbf{H}}} \equiv p_{\mathbf{H}}$. The Itakura-Saito distance, also called relative Burg entropy, is a special case of a Bregman divergence [58] and is closely related to the Kullback-Leibler divergence between two Gaussian stochastic processes [59]–[61]. To ensure that D_{IS} is well-defined, we modify the quotients in (4.5) to be $(p_{\widehat{\mathbf{H}}} + \varepsilon) / (p_{\mathbf{H}} + \varepsilon)$ with $\varepsilon \ll 1$. It is clear from the definition that D_{IS} is not symmetric with respect to its argument and hence is not a distance metric in the strict sense of the word [57]. For a dataset of K realisations of the downlink channel $\{\mathbf{H}^{(k)}\}_{k=1}^K$, we take the average of the Itakura-Saito distance, $\mathbb{E}[D_{\text{IS}}(p_{\widehat{\mathbf{H}}}, p_{\mathbf{H}})]$, similar to MSE and NMSE.

With this background on the reconstruction problem, CS-based methods and the Itakura-Saito distance argued as a suitable performance metric, we can now turn our attention to the AI-based reconstruction model itself.

4.5 A PINN reconstruction model for DL channel reconstruction

We here present a PINN-based reconstruction algorithm for the sparsely sampled DL channel estimation problem, which is based on the known physical properties of the DL channel. Specifically, we take as inspiration from CS theory the notion of sparsity in the angular domain of the transmitter antenna ports. The spectral representation of DL channel \mathbf{H}_{bs} (4.4) is a discretisation of the continuous half-sphere of departure angles from the base station antenna, as illustrated in Figure 4.1. Sparsity in the angular domain is to be understood as a concentration of gain at a small number of discrete angles of departure, representing both line-of-sight (LOS) and non-line-of-sight (NLOS) paths. These localised features on the continuous half-sphere manifold have an inherent translational symmetry, i.e., invariance to translations (the angular movement of UEs). In the discretised case shown in Figure 4.1, the projection of \mathbf{H}_{bs} onto a Cartesian grid of angles skews this symmetry, but the skewness is lower in the xz -plane than at the poles. If antennas are correctly oriented for the field of view, a UE may be presumed to have a higher probability of signal gain close to the xz -plane. This is especially true when using antenna virtualisation, i.e., when antennas elements represent beamformed antenna arrays. This should reduce the impact of the skewness of the projection and hence increase the translational symmetry. The subbands is considered a third dimension of the DL channel, such that the channel is a volumetric data structure with two spatial dimensions and one frequency dimension. Furthermore, the remaining dimensions, i.e., the RX ports N_r and

4. Neural network-based reconstruction of sparsely sounded CSI-RS

polarisation N_p , are treated as features similar to color channels in images, so that the total number of features per channel sample is $N_r N_p$.

An ANN architecture that is informed with respect to the underlying physical properties of the beamspace DL channel should preserve the physical properties of the underlying system. The rise of deep ANNs in computational imaging has a fundamental connection to symmetry; the state-of-the-art convolutional neural networks (CNNs) are equivariant to discrete translations (without max-pooling) [62]. The max-pooling regularisation, i.e., choosing the maximum of the convolutional kernel, is partly invariant to translation, but does generally not result in strong invariance [63]. A type of ANN that is well-suited for the reconstruction problem as being a regression task is the U-Net architecture [64], which is a CNN-based architecture that has shown to be effective for image segmentation tasks and medical imaging of volumetric data. U-Net is characterised by an encoder-like contracting path and a decoder-like expansive path that are more or less symmetric in structure, resulting in a U-shaped architecture and hence its name. Contrary to standard encoder-decoder architectures, the U-Net architecture has skip connections between the encoder and decoder paths, effectively increasing the number of latent spaces and encouraging better gradient flow for backpropagation. This regularisation technique, also known as residual learning, is an important ingredient in modern neural network architectures that has shown to reduce training time and improve accuracy [65].

We propose *ReconstructionUNet*, a physics-informed 3D U-Net for the reconstruction of the sparsely sampled DL channel. An overview of the reconstruction model architecture is shown in Figure 4.2. The U-Net consists of an encoder-decoder architecture with skip connections between the encoder and decoder blocks, as represented by the arrows. The input is sequentially transformed as indicated by the flow diagram in the figure. First, the (sparsely sampled) DL channel as measured by the UE is transformed to the angular domain using a DFT over the TX antenna port dimensions using (4.4). The normalisation is performed by dividing the Fourier-transformed DL channel \mathbf{H}_{bs} with the maximum

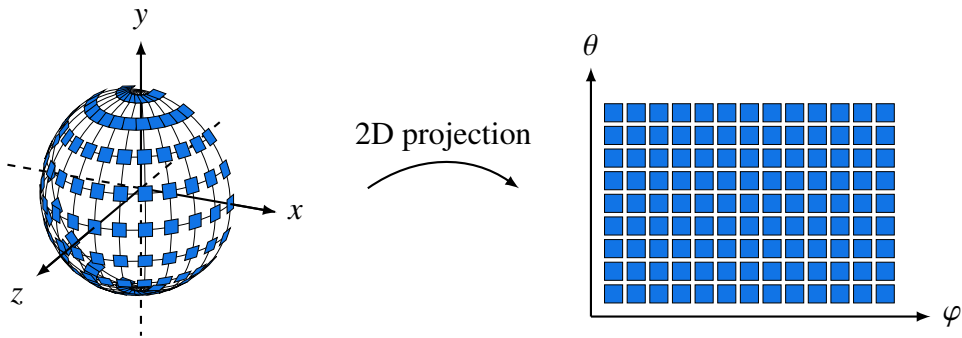


Figure 4.1: Half-sphere illustration of the DL channel in TX beamspace discretised over the transmitter antenna ports, \mathbf{H}_{bs} , and its 2D projection onto the $\varphi\theta$ -plane. Every blue square represents an entry of \mathbf{H}_{bs} at a specific AoD from the base station antenna, for a given receiver antenna port and PRB.

4. Neural network-based reconstruction of sparsely sounded CSI-RS




Figure 4.2: ReconstructionUNet architecture overview for the sparse DL channel estimation.



Figure 4.3: Convolutional block (ConvBlock) of the U-Net architecture.

absolute value of the channel sample, i.e.,

$$\text{MaxNorm}(\mathbf{H}_{\text{bs}}) := \frac{1}{a} \mathbf{H}_{\text{bs}}, \quad a = \max_i |\text{vec}(\mathbf{H}_{\text{bs}})_i|, \quad (4.6)$$

effectively scaling distribution of elements of the Fourier-transformed DL channel to fit inside the unit circle in the complex plane.

The fundamental building block of the U-Net architecture is the convolutional block (ConvBlock), which consists of two convolutional layers, each followed by a batch normalisation layer and a complex ReLU activation function, shown in Figure 4.3. The convolutional block consists of two 3D convolutional layers, each followed by a 3D batch normalisation layer and a complex-valued ReLU activation function. The output of the second convolutional layer is followed by a dropout layer. The convolutional layer (Conv3D) applies 3D convolutions to a batch of N complex-valued inputs $X \in \mathbb{C}^{N \times C_{\text{in}} \times D_{\text{in}} \times H_{\text{in}} \times W_{\text{in}}}$, where C_{in} is the number of input channels and $(D_{\text{in}}, H_{\text{in}}, W_{\text{in}})$ are the dimensions of the input tensor. We learn a set of C_{out} complex-valued kernels $\mathbf{W} \in \mathbb{C}^{C_{\text{out}} \times C_{\text{in}} \times D_{\text{out}} \times H_{\text{out}} \times W_{\text{out}}}$ and biases $\boldsymbol{\theta} \in \mathbb{C}^{C_{\text{out}}}$. For each sample $n \in \{1, \dots, N\}$ in the input batch and every output channel $o \in \{1, \dots, C_{\text{out}}\}$, the output of the convolution is given by

$$Y_{no} = \theta_o + \sum_{i=1}^{C_{\text{in}}} W_{oi} \star X_{ni}, \quad (4.7)$$

where \star is the three-dimensional cross-correlation operator [66]. To increase learning capabilities, the output of the 3D convolutional layer is normalised across the batch (BatchNorm3D in Figure 4.3) [67]. Trabelsi *et al.* [68] first extended the batch normalisation layer to complex-valued networks by viewing the problem as a two-dimensional problem of whitening, using the conversion $z = x + jy \in \mathbb{C}^N \rightarrow \tilde{z} = [x \ y]^\top \in \mathbb{R}^{2 \times N}$. The normalisation is then computed as

$$\tilde{z}_{\text{norm}} = \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}} (\tilde{z} - \hat{\boldsymbol{\mu}}), \quad (4.8)$$

where $\hat{\boldsymbol{\mu}}$ is the estimated mean $\mathbb{E}[\tilde{z}]$ and $\hat{\boldsymbol{\Sigma}}$ is the estimated covariance matrix

$$\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} \\ \sigma_{yx} & \sigma_{yy} \end{bmatrix} = \begin{bmatrix} \text{Cov}(\Re(z), \Re(z)) & \text{Cov}(\Re(z), \Im(z)) \\ \text{Cov}(\Im(z), \Re(z)) & \text{Cov}(\Im(z), \Im(z)) \end{bmatrix}, \quad (4.9)$$

4. Neural network-based reconstruction of sparsely sounded CSI-RS

with $\text{Cov}(\cdot, \cdot)$ being the covariance and $\Re(\cdot)$ and $\Im(\cdot)$ the real and imaginary parts of the complex-valued argument, respectively. During backpropagation, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are updated using momentum [69]. The decorrelation of the real and imaginary parts as is done here has proven to be advantageous in reducing overfitting of the model in the training phase [68]. The activation function in ConvBlock is an extension of the real-valued rectified linear unit (ReLU) activation function $\text{ReLU}(x) := \max(x, 0)$ to complex numbers, defined to be [68]

$$\mathbb{C}\text{ReLU}(x) := \text{ReLU}(\Re(x)) + j\text{ReLU}(\Im(x)). \quad (4.10)$$

To further prevent overfitting, a dropout layer [70] is placed as the final layer of each ConvBlock that sets a pre-determined fraction of activations to zero at random during training. This regularisation technique also improves the generalisation of the model [71].

The 3D U-Net employs two encoder blocks for the contracting path of the ANN, where the i th block is defined as the learnable operator $f_{E_i}(\cdot; \boldsymbol{\Theta}_{E_i})$ with parameters $\boldsymbol{\Theta}_{E_i}$ that transforms a batched input tensor $\mathbf{X} \in \mathbb{C}^{N \times C_{\text{in}} \times D \times H \times W}$ to a tensor $\mathbf{Y} \in \mathbb{C}^{N \times C_{\text{out}} \times D_{\text{out}} \times H_{\text{out}} \times W_{\text{out}}}$ by the computation

$$\mathbf{Y}_{E_i} = f_{E_i}(\mathbf{X}; \boldsymbol{\Theta}_{E_i}) := \text{ConvBlock} \circ \text{ConvBlock} \circ \text{MaxPool3D}(\mathbf{X}). \quad (4.11)$$

The MaxPool3D operation is here a complex-valued 3D max pooling layer that reduces the spatial dimensions of the input tensor by a factor of two. The complex-valued 3D max pooling implemented for the U-Net architecture in question is analogously to $\mathbb{C}\text{ReLU}$ a decoupled ordinary, real-valued 3D max pooling applied to the real and imaginary parts. The decoder blocks are similarly structured, where the i th block consists of an operator $f_{D_i}(\cdot; \boldsymbol{\Theta}_{D_i})$ with learnable parameters $\boldsymbol{\Theta}_{D_i}$ such that

$$\mathbf{Y}_{D_i} = f_{D_i}(\mathbf{X}; \boldsymbol{\Theta}_{D_i}) := \text{ConvTranspose3D}(\text{ConvBlock} \circ \text{ConvBlock}(\mathbf{X}) + \mathbf{Y}_{E_i}). \quad (4.12)$$

The ConvTranspose3D operator is a complex-valued 3D transposed convolutional layer that increases the spatial dimensions of the input tensor by a factor of two, also termed up-sampling, and may be viewed as the gradient of Conv3D with regards to the input [72]. The skip connections between the encoder and decoder blocks are realised by adding the output of the corresponding encoder block together with the input to the ConvTranspose3D layer. The bottleneck of the U-Net can be represented by the learnable operator $f_B(\cdot; \boldsymbol{\Theta}_B) := \text{ConvBlock} \circ \text{ConvBlock}$. The final layer of the ANN architecture is a Conv3D layer that results in the same input-output size of the ANN model. The output of the ANN model is finally scaled back using the inverse operation of MaxNorm and then then transformed back to element space using the inverse DFT.

With the theoretical methods developed for both the sparse sampling problem and the reconstruction problem with application to sparse CSI-RS sounding, the next chapter will present the experimental setup and results of the proposed methods.

5

Experiments and results

The subsequent sections describe the experimental methods by which sparse sampling and reconstruction of DL channels for a 3GPP urban microcell (UMi) scenario has been investigated, and the results obtained. First, we outline the experimental setup for the system simulation and data handling. A number of fixed spatial pilot sampling patterns are then presented, used in the experiments for comparison of different sampling strategies. The resulting near-optimal sampling pattern found by Algorithm 3.3 is then presented and discussed qualitatively and quantitatively in relation to the reference sampling patterns. Finally, the ANN-training setup and the experimental results for the ANN-based reconstruction presented and analysed in relation to the BPDN benchmark.

5.1 Simulation of DL channels and data partitioning

The experimental validation of the methods developed in Chapters 3 and 4 is here based on spatial channel model simulations in a realistic MIMO RAN setting with NLOS propagation. As part of standardisation efforts by the organisation partners of 3GPP, channel modelling at both link level and system level for MIMO simulations has been developed, including specifications for physical parameters, methodology and antenna arrangements [73]. Three main types of environments are distinguished: suburban macrocell, urban macrocell and urban microcell, the latter abbreviated as UMi. In this work, we have chosen to consider the latter environment for the simulation of DL channels and choose a carrier frequency of 2 GHz with a 30 kHz spacing of the subbands. The UMi scenario is generally defined for the case of a BS inter-site distance of less than 1 km, with antennas placed in rooftop height. For efficient coverage from available BS sites, the deployment area is divided into identical hexagonal cells as illustrated in Figure 5.1. At




Figure 5.1: Hexagonal cell layout for BS sites (orange points). Every site has three arrows corresponding to a tri-antenna setup, each serving UEs in the corresponding cell.

5. Experiments and results

Table 5.1: Simulation parameters for the UMi simulation scenario.

Parameter	Value
TX antenna array	Dual-polarised UPA, $(4 \times 1)_{SA} (8 \times 8)$, $N_t = 128$
RX antenna array	Dual-polarised UPA, $N_r = 2$
Antenna element spacing	$0,5\lambda$
Inter-site distance	200 m
Minimal distance between BS and UE	10 m
BS site elevation	10 m
UE elevation	1,5 m
Frequency	2 GHz
Subcarrier spacing	30 kHz

each BS site, three antenna arrays are placed with 120° angular spacing such that each antenna serves UEs in one hexagonal cell. With this scenario setup, we further specify the configuration used for the simulation of DL channels in Table 5.1. The antennas used are a dual-polarised UPAs that are critically spaced ($\lambda/2$) in both horizontal and vertical directions, with a total of $N_t = 8 \times 8 \times 2 = 128$ TX antenna ports with 4×1 SAs, and $N_r = 1 \times 1 \times 2$ RX antenna ports. As explained in Chapter 2, we consider a virtualised antenna array that for the UMi scenario is formed by four vertically stacked physical antenna elements.

The simulation of downlink channels was performed for 100 000 UEs randomly initialised on the hexagonal grid. With a probability of 80 %, a UE was assigned physical parameters corresponding to indoor conditions; otherwise, outdoor conditions were given to UEs. One downlink channel realisation was logged for every UE at a single time instance. It is important to note that since we only consider the noiseless case, the simulation effectively corresponds to a link-level simulation. Examples of TX radiation patterns for DL channel realisations are shown in Figure 5.2. As observed in the example radiation patterns from the dataset, there are a few dominant AoDs, with high directivity as expected for the SA beamforming. This confirms the assumption made in Chapter 4 of TX beamspace sparsity for the UMi scenario.

The dataset of logged downlink channel samples was then divided into training, validation, and test sets for unbiased evaluation of the methods developed in this thesis, as shown in Figure 5.3. The training set consisting of 80 % of samples was further partitioned in half and used for two purposes: Pre-training the ANN model detailed in Section 4.1 and finding near-optimal sampling patterns for antenna port muting using submodular optimisation, and secondly to fine-tune the ANN model to specific antenna port mutings. To optimise hyperparameters of the ANN model, a 20 % validation set was used in both training phases to evaluate the performance of the model. Finally, the test set of 20 000 DL channel realisations was used to evaluate the performance of the trained model on unseen data, and compare the performance to the BPDN algorithm.




Figure 5.2: TX radiation patterns for the UMi simulation. Each subplot represents a realisation of the downlink channel for a UE at a certain time instance, and for a single receiver antenna port and PRB. The channel is shown as a heatmap of relative spectral power gain for zenith (θ) and azimuth (ϕ) departure angles, with a maximum gain of zero.

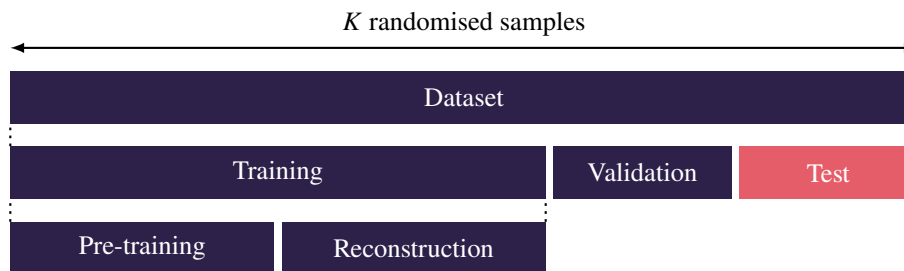



Figure 5.3: Data partitioning for the simulated downlink channels. The training set is further split into two parts; one part used for pre-training and the other for fine-tuning reconstruction with specific antenna port mutings of the ANN model.


5.2 Fixed spatial sampling patterns for CSI-RS

For the reconstruction of the full DL channel from the sparse measurements, we consider a number of fixed spatial sampling patterns for antenna port muting, presented in Figures 5.4 to 5.6. The chosen muting patterns consists of combinations 25 % and 50 % muting levels, (uniformly) random and regular (‘checkered’) patterns, as well as joint and separate muting across the two polarisations. The random muting patterns were fixated before the experiments by a single random draw and used for all experiments involving random muting. Although computing averages over multiple randomly drawn muting patterns is preferable for a better understanding of the full distribution of patterns, limitations of computational resources for the training of the ANN model made this approach unfeasible. To aid the reader in comparing radiation patterns for different antenna port sounding in the following sections, these fixed patterns are presented again on the side in such plots.

5. Experiments and results




(a) 25 % uniformly random muting.




(b) 50 % uniformly random muting.

Figure 5.4: Randomised spatial antenna port sampling patterns with for the $N_t = 8 \times 8 \times 2 = 128$ port UPA, for (a) 25 % and (b) 50 % muting levels.




(a) 25 % uniformly random muting
(jointly).




(b) 50 % uniformly random muting
(jointly).

Figure 5.5: Randomised spatial antenna port sampling patterns with for the $N_t = 8 \times 8 \times 2 = 128$ port UPA, for (a) 25 % and (b) 50 % muting levels, jointly across polarisations.



(a) 50 % checkered muting.



(b) 50 % checkered muting (jointly).

Figure 5.6: Checkered spatial antenna port sampling patterns with for the $N_t = 8 \times 8 \times 2 = 128$ port UPA, for (a) 50 % muting level and (b) 50 % muting level jointly across polarisations.

5.3 Near-optimal spatial sampling patterns for CSI-RS

To extend the set of fixed spatial sampling patterns for CSI-RS antenna ports, we here present the results of the submodular optimisation of spatial sampling patterns for the $N_t = 8 \times 8 \times 2 = 128$ port UPA using Algorithm 3.3 with average of G (3.18) as the objective function (i.e., $\mathbb{E}_w[G_p^{(k)}]$ with $w_k = 1$ for all k and $p = 1$). The greedy algorithm was run on a single NVIDIA GeForce RTX 3090 GPU with a total computation time of less than 15 minutes for all N_t antenna ports and the $K = 30\,000$ channel realisation corresponding to the pre-training dataset. The upper and lower bounds on the submodular function values were also computed by greedily choosing the antenna port that maximises $G^{(k)}$ across all k data points for the upper bound, $\max_k G^{(k)}$, and the antenna port that minimises $G^{(k)}$ for the lower bound, $\min_k G^{(k)}$; by iteratively selecting the element that yields the smallest/highest possible value of $G^{(k)}$ across the dataset, bounds representative of the values of $G^{(k)}$ are created, with an identical computational cost as Algorithm 3.3. The computed maximal average submodular function values for the greedy algorithm are shown in Figure 5.7. The coloured areas represents the bounds of $G^{(k)}$. The plotted red line shows representing the maximal average submodular function found by the greedy algorithm is closer to the upper bound than the lower bound but relatively close to the mean of the two bounds. Figuratively speaking, Algorithm 3.3 ‘draws’ the red line from left to right in Figure 5.7 as more antenna ports are removed and \mathcal{S} grows, until in this case $L = |\mathcal{S} \setminus \mathcal{N}| = 0$. An important distinction to make from this result is that submodularity necessarily does not hold for $\max_k G^{(k)}$ and $\min_k G^{(k)}$, from which the bounds are derived. Although the maximal average submodular function is approximately concave by visual inspection, submodularity does not imply concavity in the graph of




Figure 5.7: Computed submodular function values for Algorithm 3.3 with upper and lower bounds across the test dataset, with an $N_t = 8 \times 8 \times 2 = 128$ port UPA and separate muting for polarisations.

5. Experiments and results

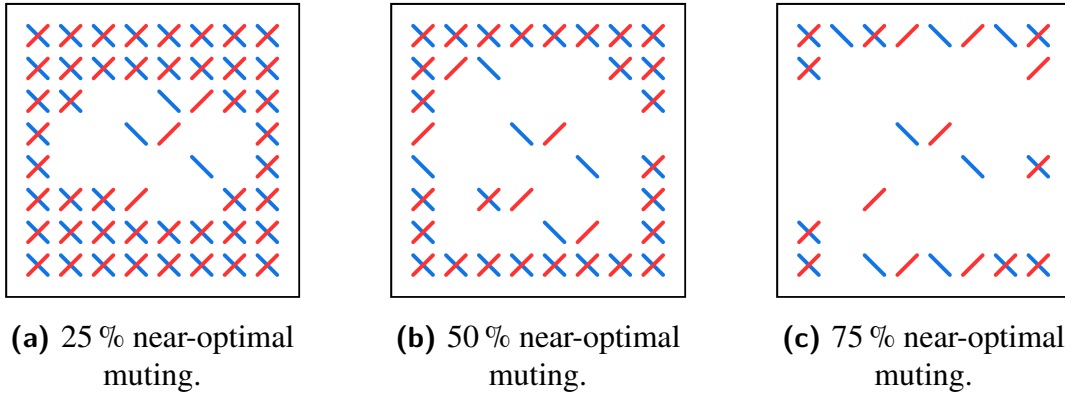


Figure 5.8: Near-optimal spatial antenna port sampling patterns from submodular optimisation using Algorithm 3.3 for the $N_t = 8 \times 8 \times 2 = 128$ port UPA, here shown for (a) 25 %, (b) 50 % and (b) 75 % muting levels. The sampling patterns has average submodular function values of 0,455, 0,754 and 0,938, respectively.

the submodular function. According to the definition, the property of diminishing returns applies to specific elements added to the set \mathcal{S} , not to the case of adding an arbitrary element to \mathcal{S} , and hence no concavity is guaranteed.

By running Algorithm 3.3 once for $L = 0$, near-optimal spatial sampling patterns were found for all possible number muted ports $1, \dots, N_t - 1$. Figure 5.8 shows the found near-optimal antenna port sampling patterns for the 25 %, 50 % and 75 % muting levels. Evidently, the found near-optimal spatial sampling patterns have a clear spatial structure in that the distance between the sampled antenna ports in general is maximised, and secondly that the sampling is denser in vertical rather than horizontal direction. The greedy algorithm effectively removes antenna ports that are located in the centre of the array, with an emphasis on the horizontal direction. It is notable that the found near-optimal spatial sampling pattern for 50 % muting level is at least 75 %-optimal with respect to the submodular function G , as the average submodular function value is 0,754. Similarly, the found near-optimal spatial sampling pattern for 75 % muting level is at least 93 %-optimal with respect to G .

To place the near-optimal sampling patterns into context, we use the Itakura-Saito distance (4.5) with $\varepsilon = 10^{-15}$ for the sparsely sampled DL channel PSD p_{H_L} . Figure 5.9 shows the average Itakura-Saito Distances for the near-optimal and fixed muting patterns antenna port sampling patterns for the 50 % muting level, computed across the test dataset. It is clear from the figure that the near-optimal muting pattern results in the Itakura-Saito distance by at least a factor of two as compared to the checkered sampling patterns. In the 25 % muting case, the Itakura-Saito distances for the three antenna port sampling patterns are shown in Figure 5.10. As expected, the distances are lower for all respective types of sampling patterns when fewer antennas are muted, and in particular the near-optimal sampling pattern results in the shortest overall distance. In fact, the farthest distance in the 25 % case is lower than the shortest distance in the 50 % case.

A qualitative complement to the Itakura-Saito distance is the visualisation of the channel power spectrum for the near-optimal muting patterns, in comparison to the fixed muting patterns. Figure 5.11 shows an example of a radiation pattern for a DL channel




Figure 5.9: Itakura-Saito distance for the near-optimal muting patterns with 50% muting level, as compared to the fixed sampling patterns.




Figure 5.10: Itakura-Saito distance for the near-optimal muting patterns with 25% muting level, as compared to the fixed random sampling patterns.

realisation with 50% muting level. The radiation patterns shown are analogous to those in Figure 5.2, except that the axes have been omitted. The random muting patterns and the near-optimal patterns preserve the peak gain angles although there is visible noise, whereas the checkered sampling patterns in Figures 5.11e and 5.11f visibly exhibit less noise but instead aliasing effects. One possible reason for the aliasing effects is the loss of inter-antenna port distance resolution. By visual inspection, the radiation patterns for all sampling patterns with 25% muting level are similar in relative angular power gain and show a lower noise than as compared to the corresponding patterns with in the 50% muting case.

5. Experiments and results

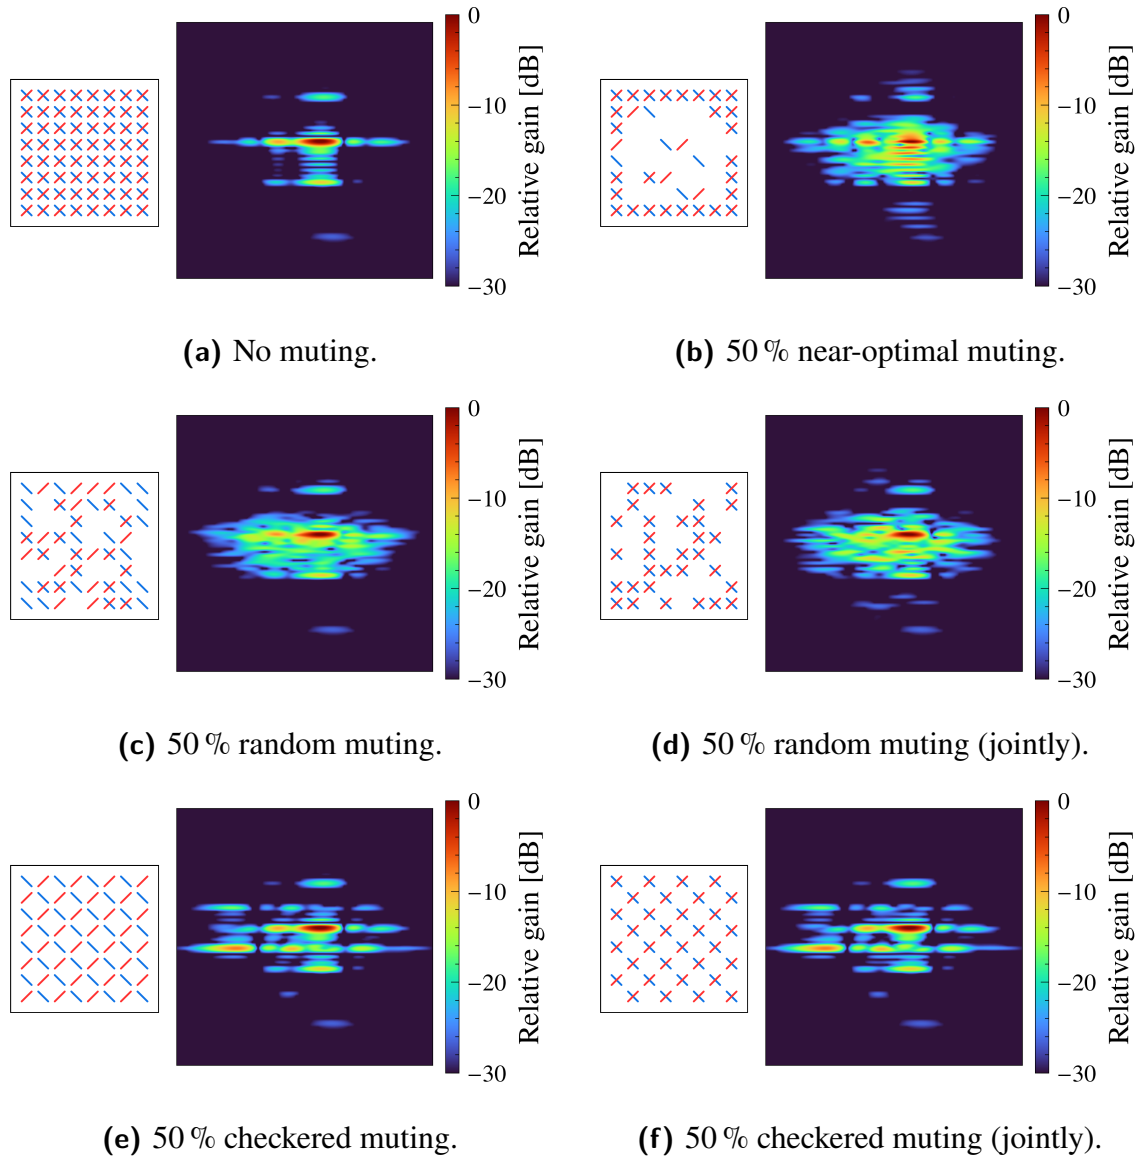


Figure 5.11: TX radiation patterns for the fixed UPA sampling patterns for the 25% and 50% muting levels. The left column shows the antenna port sampling pattern and the right column shows the corresponding channel power spectrum.

5. Experiments and results

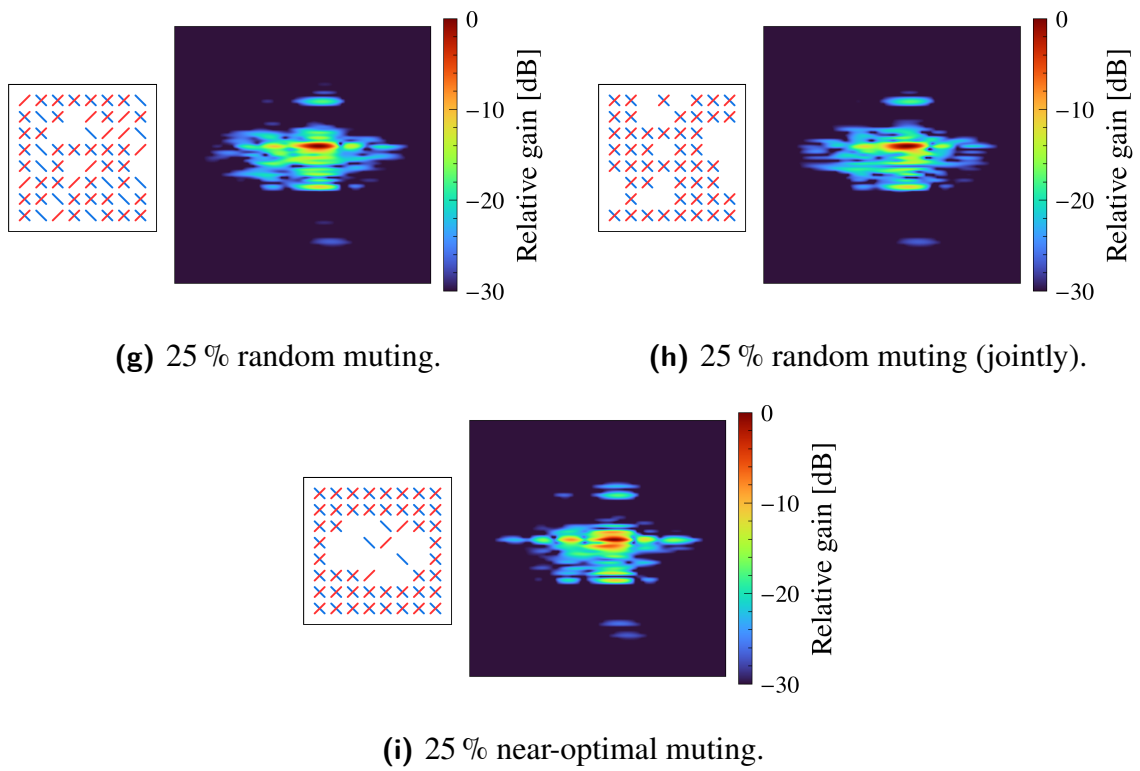


Figure 5.11: TX radiation patterns for the fixed UPA sampling patterns for the 25 % and 50 % muting levels (cont.).

5. Experiments and results

5.4 Neural network-based reconstruction of the DL channel for CSI-RS

This section presents the training procedure and the results of the ANN-based reconstruction of the DL channel from sparse CSI-RS measurements using the ReconstructionUNet model detailed in Section 4.5, with qualitative and quantitative comparisons of reconstruction accuracy to the BPDN algorithm, and of computational performance.

5.4.1 Pre-training and fine-tuning of the neural network

The learning task for an ANN model of accurately reconstructing the estimated DL channel from CSI-RSs is a challenging task in the sense that the model has to both learn to efficiently represent partial DL channels on a low-dimensional manifold and to recover the full DL channel tensor from a sparse sampling scheme. From this perspective it is pertinent to train the ANN model to solve problems of incremental difficulty; this way, the model first learns the representation of the physical system in a pre-training phase, and is in a second phase fine-tuned with lower computational effort for specific scenarios. In the regression task at hand, we therefore first let the ReconstructionUNet learn to accurately reproduce the input and then fine-tune the model for specific antenna port sampling patterns. A number of other well-known techniques was also used to maximise the learning outcome. MSE was selected as the loss function for backpropagation for its simplicity and interpretability. The learning rate at which backpropagation errors were properly set by an initial sweep of the loss function for 30 different learning rates in the range 10^{-8} – 10^3 at the start of each training phase, ensuring that the learning rate is adjusted to the errors computed in the MSE. The adaptive moment estimation (Adam) optimisation algorithm [74] has become one of the most frequently used optimisers in deep learning [75] and is used for computing the gradients in the backpropagation of the ANN model. To prevent overfitting, the loss of the validation dataset was monitored and training concluded when no improvement of the loss was observed for a number of training epochs. The ReconstructionUNet model was trained on an NVIDIA GeForce RTX 3090 GPU for 15 h in the pre-training phase. Fine-tuning the model to the different spatial sampling patterns took between 5 h and 12 h.

5.4.2 Reconstruction of DL channels

We here present the results of ReconstructionUNet and compare its performance to the BPDN algorithm (O4.7). The threshold of BPDN was set sufficiently small; that is, the threshold was set to be at least one order of magnitude smaller than the reconstruction performance (in this case 10^{-4}). As a first point of comparison, Table 5.2 shows the average computational performance of the two reconstruction methods. The measurement was done in terms of number of floating point operations (FLOP), inference speed (FLOP/s) and latency, averaged over the test dataset for the fixed uniformly random spatial sampling pattern with 50 % muting level (see Figure 5.4b). Although the computational effort to reconstruct a sparsely sampled DL channel is larger for the ANN model, the optimisation

Table 5.2: Average computational performance of ReconstructionUNet and BPDN per downlink channel.

Reconstruction method	GFLOP	GFLOP/s	Latency
ReconstructionUNet	2,07	40,66	51 ms
BPDN	0,48	0,03	15 s

Table 5.3: Reconstruction performance of ReconstructionUNet and BPDN in terms of NMSE, with 25 % and 50 % muting antenna port sampling patterns.

Antenna port sampling pattern	NMSE [dB]		Muting level
	ReconstructionUNet	BPDN	
Checkered	-4,04	-1,53	50 %
Checkered (jointly)	-2,57	-1,45	
Random	-3,60	-5,67	
Random (jointly)	-3,41	-4,81	
Near-optimal	-3,41	-4,46	
Random	-4,03	-11,23	25 %
Random (jointly)	-3,82	-9,67	
Near-optimal	-3,92	-9,86	

algorithm performed in BPDN is significantly slower, resulting in a latency of 15 s for a single DL channel reconstruction as compared to 51 ms for the ReconstructionUNet. Note that the ANN model computation was accelerated on a GPU, whereas the BPDN model was run on a single CPU core. However, even if run on a CPU, the ReconstructionUNet outperforms the BPDN algorithm in inference speed by an order of magnitude.

The reconstruction performance of the ReconstructUNet model and BPDN algorithm are shown in Table 5.3. We observe that BPDN in all cases except for the regular patterns has a lower average NMSE reconstruction error than the ANN model. In particular for the 25 % muting case, the difference is significant; the ReconstructionUNet here shows a similar performance as to for the reconstruction of patterns with 50 % muting level. The antenna port sampling pattern that yields smallest reconstruction error for BPDN is the random pattern, for both muting levels. For the ANN model on the other hand, the checkered sampling pattern results most accurately reconstruction in terms of NMSE for the 50 % muting level, whereas for 50 % muting the random pattern shows the lowest error. The worst sampling pattern for both reconstruction algorithms is the checkered sampling pattern with joint muting across polarisations.

The other performance metric evaluated for the reconstruction methods is the Itakura-Saito distance, shown for the 50 % and 25 % muting levels in Figures 5.12 and 5.13, respectively. The ReconstructionUNet here outperforms BPDN for all sampling patterns and muting levels. For the 50 % muting level, the Itakura-Saito distance for the near-optimal sampling pattern is the shortest for both methods, with a factor of eight shorter distance for the ANN model as compared to BPDN. At the 25 % muting, the

5. Experiments and results




Figure 5.12: Reconstruction performance of ReconstructionUNet and BPDN in terms of the average Itakura-Saito distance, $\mathbb{E}[D_{\text{IS}}(p_{\hat{H}}, p_H)]$, with 50 % muting antenna port sampling patterns.




Figure 5.13: Reconstruction performance of ReconstructionUNet and BPDN in terms of the average Itakura-Saito distance, $\mathbb{E}[D_{\text{IS}}(p_{\hat{H}}, p_H)]$, with 25 % muting antenna port sampling patterns.

jointly muted random sampling pattern yields the shortest Itakura-Saito distance for the ANN model, whereas the separately muted random sampling pattern results in the shortest distance for BPDN. Notably, all sampling patterns with 25 % muting level have a significantly shorter Itakura-Saito distance than the corresponding sampling patterns with 50 % muting level for the two reconstruction methods.

To visualise the reconstructed DL channel, Figure 5.14 shows the TX radiation patterns of the same DL channel realisation, receiver antenna and PRB as in Figure 5.11, but here for the recovered DL channel using the ReconstructionUNet and BPDN. Although only one example of a channel reconstruction for on RX port and one subband is shown, the radiation patterns here compound the results of the computed Itakura-Saito distances. Visibly, for 50 % muting, the radiation pattern of the reconstructed channel has less noise and less aliasing with ReconstructionUNet versus with the BPDN algorithm. For 25 % muting on the other hand, the channel reconstructions for different sampling patterns and for the two reconstruction methods are indistinguishable in their radiation patterns for this particular channel realisation. It is noticeable in the comparison of the NMSE and the Itakura-Saito distance results is that there is a significant difference. There are many possible reasons for this, but the author notes that improvements to the implementation of the ANN are likely to yield a significant increase in reconstruction accuracy in terms of the NMSE, thereby qualitatively aligning to a greater extent to the displayed Itakura-Saito distances.

5. Experiments and results

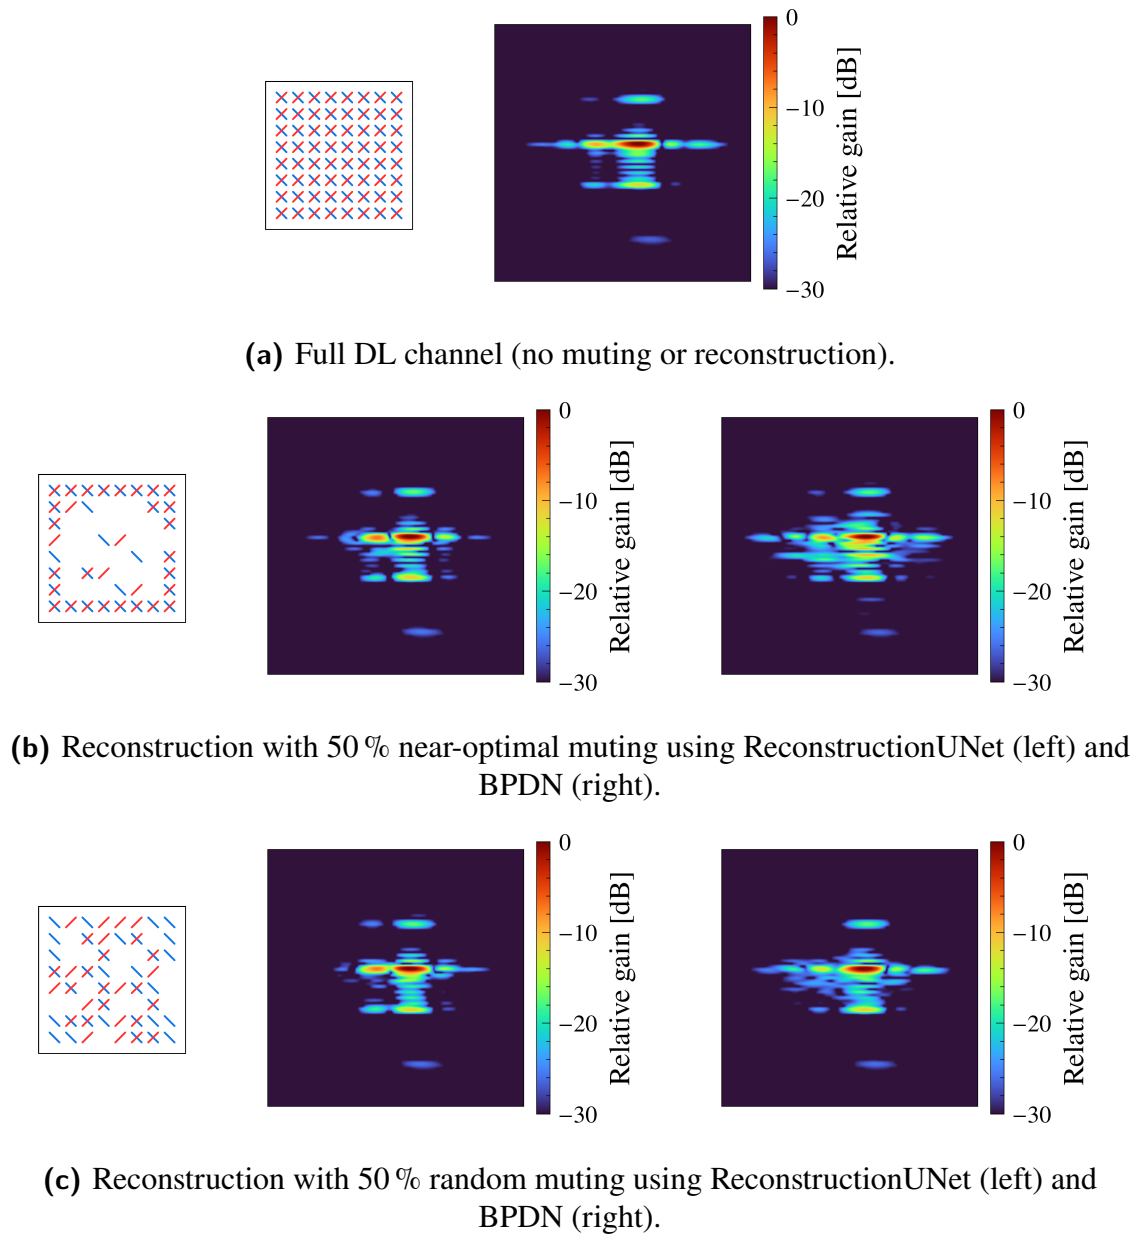




Figure 5.14: TX radiation patterns for the reconstruction of a DL channel realisation with fixed UPA sampling patterns for 25% and 50% muting levels using ReconstructionUNet and BPDN.


5. Experiments and results



(d) Reconstruction with 50 % random muting (jointly) using ReconstructionUNet (left) and BPDN (right).




(e) Reconstruction with 50 % checked muting using ReconstructionUNet (left) and BPDN (right).




(f) Reconstruction with 50 % checked muting (jointly) using ReconstructionUNet (left) and BPDN (right).


Figure 5.14: TX radiation patterns for the reconstruction of a DL channel realisation with fixed UPA sampling patterns for 25 % and 50 % muting levels muting levels using ReconstructionUNet and BPDN (cont.).



(g) Reconstruction with 25 % random muting using ReconstructionUNet (left) and BPDN (right).



(h) Reconstruction with 25 % random muting (jointly) using ReconstructionUNet (left) and BPDN (right).



(i) Reconstruction with 25 % near-optimal muting using ReconstructionUNet (left) and BPDN (right).

Figure 5.14: TX radiation patterns for the reconstruction of a DL channel realisation with fixed UPA sampling patterns for 25 % and 50 % muting levels muting levels using ReconstructionUNet and BPDN (cont.).

5. Experiments and results

6

Conclusion

In this thesis, we have investigated the problem of sparse sampling and reconstruction of DL channels for feedback-based CSI acquisition in massive MIMO systems. Our contributions are mainly methodological, although experimental results are also presented. We have proposed a twofold approach to reduce the overhead of DL CSI-RS in massive MIMO systems: optimising pilot placement for the sparse sounding of the reference signals, and reconstructing the full channel information from these sparse measurements at the UE.

The sparse pilot placement problem was formulated as a submodular optimisation problem, where we defined a cost function based on the p -frame potential of the DL channel. This proposed objective function allowed us to find near-optimal solutions with improved optimality guarantees for the sparse pilot placement problem by using a greedy algorithm on a dataset of simulated DL channels, both with respect to the cost function and to the p -frame potential. Based on simulated data in a 3GPP UMi environment with UPAs deployed at the BS and UE and noiseless conditions, we show that the found near-optimal sampling patterns for the pilots retain the PSD of the DL channel in terms of the Itakura-Saito distance measure, while reducing the overhead associated with the transmission of CSI-RS to a greater extent than the fixed pilot placements considered in this work.

We also have proposed ReconstructionUNet as a PINN model for the reconstruction of the full DL channel from the estimated partial DL channel from spatially sparse CSI-RS measurements. The U-Net architecture is designed to leverage the sparse angular representation of the DL channel, allowing for efficient reconstruction of the full channel information. The U-Net was trained in a two-stage process on a large dataset of simulated noiseless DL channels for the same 3GPP environment and for several different spatial pilot configurations and muting levels. The performance of the U-Net was evaluated in terms of reconstruction accuracy and computational efficiency with respect to inference latency, and compared to the performance of BPDN as a classical CS-based method. The results show that the U-Net model achieves a comparable reconstruction performance in NMSE and the Itakura-Saito distance measure as compared to BPDN. The ANN model also outperforms BPDN in terms of computational efficiency with respect to floating point operations per second and overall latency in the experimental settings, demonstrating the potential of AI-based techniques for improving the performance of massive MIMO systems in upcoming 6G networks and beyond. The performance of ReconstructionUNet and the BPDN algorithm was evaluated for different pilot selection strategies, including random and fixed pilot placements. From the experiments, we conclude that the choice of pilot selection strategy had a noticeable impact on the reconstruction performance of

6. Conclusion

both methods in the considered scenario, to a greater extent in the case of a 50% pilot muting level than for 25% muting. These results indicate that the combined approach of optimising sparse pilot placement and using a neural network for reconstruction can effectively reduce the overhead of DL CSI-RS in massive MIMO systems.

6.1 Future work

To further validate the proposed near-optimal antenna port sampling patterns and the performance of ReconstructionUNet, we recommend conducting additional experiments in system-level simulations to assess the impact on overhead associated with the transmission of CSI-RS in a realistic network environment, in combination with the generalisability of AI models to other cellular scenarios, channel distributions, antenna arrays, and imperfect CSI with finite SINR. The proposed ANN architecture only preserves the translational equivariance of the MIMO channel in the TX angular domain to an approximation; fully preserving the spherical structure, for example by using spherical CNNs, would in theory allow for a lower-footprint ANN with equal performance. In addition, it is recommended to further investigate how other random antenna port sampling patterns compares to the other fixed muting patterns. Although the frame potential has a theoretical justification for its use as a cost function for the sparse sampling problem, it would be interesting to investigate other cost functions and compare their performance to the frame potential, and furthermore to better understand the resulting near-optimal sampling patterns quantitatively. The research conducted in this thesis work has not addressed the overhead reduction of CSI reporting from a standardisation point of view, and the constraints this imposes on possible solutions, which should be taken into consideration for real-world application for RAN in massive MIMO.

Bibliography

- [1] ‘Ericsson Mobility Report November 2024,’ Ericsson, Stockholm, Sweden, Tech. Rep., 2024. [Online]. Available: <https://www.ericsson.com/en/reports-and-papers/mobility-report/reports/november-2024>.
- [2] H. Asplund, D. Astely, P. von Butovitsch *et al.*, *Advanced Antenna Systems for 5G Network Deployments: Bridging the Gap Between Theory and Practice*. London, UK: Academic Press, 2020, ISBN: 9780128200469.
- [3] S. Muruganathan, F. Athley, S. Gao and X. Zhang, ‘Spatial domain CSI compression for coherent joint transmission,’ US Patent Appl. 20250055528A1, Aug. 23, 2022.
- [4] ‘3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Release 18 Description; Summary of Rel.18 Work Items (Release 18),’ 3GPP, Valbonne, France, Tech. Rep. 21.918, version 2.2.0, 2025. [Online]. Available: <http://www.3gpp.org>.
- [5] O. Liberg, C. Hoymann, C. Tidestav *et al.*, ‘Introducing 5G Advanced,’ *IEEE Commun. Stand. Mag.*, vol. 8, no. 1, pp. 52–57, 2024. DOI: 10.1109/MCOMSTD.0003.2200059.
- [6] I. J. Good, ‘The Interaction Algorithm and Practical Fourier Analysis,’ *J. R. Stat. Soc. Ser. B. Stat. Method*, vol. 20, no. 2, pp. 361–372, 1958, ISSN: 00359246. [Online]. Available: <http://www.jstor.org/stable/2983896>.
- [7] H. Asplund, D. Astely, P. von Butovitsch *et al.*, ‘Ericsson Massive MIMO Handbook: Extended version,’ Ericsson AB, Tech. Rep. EN/LZT 4/28701-FGB1010987 Uen Rev C, 2024. [Online]. Available: <https://www.ericsson.com/massive-mimo>.
- [8] E. Dahlman, S. Parkvall and J. Sköld, *5G NR: The Next Generation Wireless Access Technology*, 2nd. London, UK: Academic Press, 2020, ISBN: 9780128223215.
- [9] E. L. Bengtsson, ‘Massive MIMO From a Terminal Perspective,’ Doctoral Thesis (compilation), Department of Electrical and Information Technology, Lund University, 2019, ISBN: 978-91-7895-253-3.
- [10] X. Fu, D. Le Ruyet, R. Visoz *et al.*, ‘A Tutorial on Downlink Precoder Selection Strategies for 3GPP MIMO Codebooks,’ *IEEE Access*, vol. 11, pp. 138 897–138 922, 2023. DOI: 10.1109/ACCESS.2023.3338866.
- [11] S. P. Chepuri, ‘Sparse Sensing for Statistical Inference: Theory, Algorithms, and Applications,’ Ph.D. dissertation, TU Delft, Delft, The Netherlands, 2016. [Online]. Available: <https://theses.eurasip.org/theses/648/sparse-sensing-for-statistical-inference-theory>.

Bibliography

- [12] E. L. Lawler and D. E. Wood, 'Branch-And-Bound Methods: A Survey,' *Oper. Res.*, vol. 14, no. 4, pp. 699–719, 1966, ISSN: 15265463.
- [13] W. J. Welch, 'Branch-and-Bound Search for Experimental Designs Based on D Optimality and Other Criteria,' *Technometrics*, vol. 24, no. 1, pp. 41–48, 1982, ISSN: 00401706. DOI: 10.2307/1267576.
- [14] S. Joshi and S. Boyd, 'Sensor Selection via Convex Optimization,' *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, 2009. DOI: 10.1109/TSP.2008.2007095.
- [15] S. P. Chepuri and G. Leus, 'Sparsity-Promoting Sensor Selection for Non-Linear Measurement Models,' *IEEE Trans. Signal Process.*, vol. 63, no. 3, pp. 684–698, 2015. DOI: 10.1109/TSP.2014.2379662.
- [16] G. L. Nemhauser, L. A. Wolsey and M. L. Fisher, 'An analysis of approximations for maximizing submodular set functions—I,' *Math. Program.*, vol. 14, no. 1, pp. 265–294, 1978, ISSN: 1436-4646. DOI: 10.1007/BF01588971.
- [17] F. Bach, 'Learning with Submodular Functions: A Convex Optimization Perspective,' *Found. Trends Mach. Learn.*, vol. 6, no. 2-3, pp. 145–373, 2013. DOI: 10.1561/22000000039.
- [18] U. Feige, 'A Threshold of $\ln n$ for Approximating Set Cover,' *JACM*, vol. 45, no. 4, pp. 634–652, 1998, ISSN: 0004-5411. DOI: 10.1145/285055.285059.
- [19] J. Ranieri, A. Chebira and M. Vetterli, 'Near-Optimal Sensor Placement for Linear Inverse Problems,' *IEEE Trans. Signal Process.*, vol. 62, no. 5, pp. 1135–1146, 2014. DOI: 10.1109/TSP.2014.2299518.
- [20] L. Baldassarre, Y.-H. Li, J. Scarlett, B. Gözcü, I. Bogunovic and V. Cevher, 'Learning-Based Compressive Subsampling,' *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 809–822, 2016. DOI: 10.1109/JSTSP.2016.2548442.
- [21] R. J. Duffin and A. C. Schaeffer, 'A Class of Nonharmonic Fourier Series,' *Trans. Am. Math. Soc.*, vol. 72, no. 2, pp. 341–366, 1952. DOI: 10.2307/1990760.
- [22] J. J. Benedetto and P. J. S. G. Ferreira, Eds., *Modern Sampling Theory: Mathematics and Applications* (Applied and Numerical Harmonic Analysis). Boston, MA, USA: Birkhäuser, 2001, ISBN: 978-1-4612-0143-4. DOI: 10.1007/978-1-4612-0143-4.
- [23] J. J. Benedetto and M. Fickus, 'Finite Normalized Tight Frames,' *Adv. Comput. Math.*, vol. 18, no. 2, pp. 357–385, 2003, ISSN: 1572-9044. DOI: 10.1023/A:1021323312367.
- [24] F. Krahmer, G. Kutyniok and J. Lemvig, 'Sparse matrices in frame theory,' *Comput. Stat.*, vol. 29, no. 3, pp. 547–568, 2013. DOI: 10.1007/s00180-013-0446-1.
- [25] J. Ranieri, 'Sensing the real world: Inverse problems, sparsity and sensor placement,' Ph.D. dissertation, EPFL, Lausanne, 2014. DOI: 10.5075/epfl-thesis-6349.
- [26] M. Ehler and K. Okoudjou, 'Minimization of the probabilistic p-frame potential,' *J. Stat. Plan. Inference*, vol. 142, no. 3, pp. 645–659, 2012, ISSN: 0378-3758. DOI: 10.1016/j.jspi.2011.09.001.
- [27] J. Kovacevic and A. Chebira, 'Life Beyond Bases: The Advent of Frames (Part I),' *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 86–104, 2007. DOI: 10.1109/MSP.2007.4286567.

- [28] M. Rupf and J. L. Massey, 'Optimum sequence multisets for synchronous code-division multiple-access channels,' *IEEE Trans. Inf. Theory*, vol. 40, no. 4, pp. 1261–1266, 1994. DOI: 10.1109/18.335940.
- [29] S. Rao, S. P. Chepuri and G. Leus, 'Greedy sensor selection for non-linear models,' in *IEEE 6th Int. Workshop Comput. Adv. Multi-Sensor Adapt. Process.*, Cancun, Mexico: IEEE, 2015, pp. 241–244. DOI: 10.1109/CAMSAP.2015.7383781.
- [30] F. Zocco, M. Maggipinto, G. A. Susto and S. McLoone, 'Lazy FSCA for unsupervised variable selection,' *Eng. Appl. Artif. Intell.*, vol. 124, p. 106624, 2023, ISSN: 0952-1976. DOI: 10.1016/j.engappai.2023.106624.
- [31] G. Ortiz-Jiménez, M. Coutino, S. P. Chepuri and G. Leus, 'Sparse Sampling for Inverse Problems With Tensors,' *IEEE Trans. Signal Process.*, vol. 67, no. 12, pp. 3272–3286, 2019. DOI: 10.1109/TSP.2019.2914879.
- [32] Z. Chen, J. Ranieri, R. Zhang and M. Vetterli, 'DASS: Distributed Adaptive Sparse Sensing,' *IEEE Trans. Wireless Commun.*, vol. 14, no. 5, pp. 2571–2583, 2015. DOI: 10.1109/TWC.2014.2388232.
- [33] S. Ding, N. Tong, Y. Zhang, X. Hu and X. Zhao, 'Cognitive Antenna Selection in MIMO Imaging Radar,' *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 12, pp. 9829–9841, 2021. DOI: 10.1109/TGRS.2020.3047610.
- [34] W. Feng, P. Wan, X. Hu, Y. Guo and H. Zhu, 'Cognitive Sparse Imaging Method for MIMO Radar under Wideband Interference,' *Remote Sens.*, vol. 14, no. 21, 2022, ISSN: 2072-4292. DOI: 10.3390/rs14215294.
- [35] R. Vaze and H. Ganapathy, 'Sub-Modularity and Antenna Selection in MIMO Systems,' *IEEE Commun. Lett.*, vol. 16, no. 9, pp. 1446–1449, 2012. DOI: 10.1109/LCOMM.2012.070512.120912.
- [36] A. Konar and N. D. Sidiropoulos, 'A Simple and Effective Approach for Transmit Antenna Selection in Multiuser Massive MIMO Leveraging Submodularity,' *IEEE Trans. Signal Process.*, vol. 66, no. 18, pp. 4869–4883, 2018. DOI: 10.1109/TSP.2018.2863654.
- [37] J. Zhang, J. Wang and Y. Wang, 'Antenna selection in massive MIMO systems utilizing the submodular function,' in *9th Int. Conf. Wireless Commun. Signal Process.*, Nanjing, China: IEEE, 2018, pp. 1–6. DOI: 10.1109/WCSP.2017.8171056.
- [38] S. Muruganathan, F. Athley, S. Gao, A. Nilsson and X. Zhang, 'Frequency domain CSI compression for coherent joint transmission,' US Patent Appl. 20240364400A1, Aug. 23, 2022.
- [39] E. J. Candès, J. Romberg and T. Tao, 'Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,' *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, 2006. DOI: 10.1109/TIT.2005.862083.
- [40] D. L. Donoho, 'Compressed sensing,' *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006. DOI: 10.1109/TIT.2006.871582.
- [41] E. Crespo Marques, N. Maciel, L. Naviner, H. Cai and J. Yang, 'A Review of Sparse Recovery Algorithms,' *IEEE Access*, vol. 7, pp. 1300–1322, 2019. DOI: 10.1109/ACCESS.2018.2886471.
- [42] M. R. Garey and D. S. Johnson, *Computers and Intractability; A Guide to the Theory of NP-Completeness*. USA: W. H. Freeman & Co., 1990, ISBN: 0716710455.

Bibliography

- [43] S. G. Mallat and Z. Zhang, ‘Matching pursuits with time-frequency dictionaries,’ *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, 1993. doi: 10.1109/78.258082.
- [44] J. A. Tropp and A. C. Gilbert, ‘Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit,’ *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007. doi: 10.1109/TIT.2007.909108.
- [45] E. J. Candes, J. Romberg and T. Tao, ‘Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,’ *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, 2006. doi: 10.1109/TIT.2005.862083.
- [46] T. Piipponen, ‘Design and Analysis of Antenna Selection in Massive MIMO,’ M.S. thesis, Aalto University, Aalto, Finland, 2021. [Online]. Available: <https://urn.fi/URN:NBN:fi:aalto-202108298506>.
- [47] C. Ramirez, V. Kreinovich and M. Argaez, ‘Why ℓ_1 is a good approximation to ℓ_0 : A geometric explanation,’ *J. Uncertain Syst.*, vol. 7, no. 3, pp. 203–207, 2013.
- [48] S. S. Chen and D. L. Donoho, ‘Basis pursuit,’ in *Proc. Asilomar Conf. Signals Syst. Comput.*, vol. 1, Pacific Grove, CA, USA: IEEE, 1994, 41–44 vol.1. doi: 10.1109/ACSSC.1994.471413.
- [49] S. S. Chen, ‘Basis Pursuit,’ Ph.D. dissertation, Department of Statistics, Stanford University, Standford, USA, 1995.
- [50] D. L. Donoho and I. M. Johnstone, ‘Empirical Atomic Decomposition,’ Unpublished manuscript, 1995.
- [51] T. Georgiou and A. Lindquist, ‘Kullback-Leibler approximation of spectral density functions,’ in *42nd IEEE Int. Conf. Decis. Control*, vol. 4, IEEE, 2003, pp. 4237–4242. doi: 10.1109/CDC.2003.1271815.
- [52] M. Pavon and A. Ferrante, ‘On the Georgiou-Lindquist approach to constrained Kullback-Leibler approximation of spectral densities,’ *IEEE Trans. Autom. Control*, vol. 51, no. 4, pp. 639–644, 2006. doi: 10.1109/TAC.2006.872755.
- [53] A. Ferrante, M. Pavon and F. Ramponi, ‘Hellinger Versus Kullback-Leibler Multivariable Spectrum Approximation,’ *IEEE Trans. Autom. Control*, vol. 53, no. 4, pp. 954–967, 2008. doi: 10.1109/TAC.2008.920238.
- [54] P. Enqvist and J. Karlsson, ‘Minimal Itakura-Saito distance and covariance interpolation,’ in *47th IEEE Conf. Decis. Control*, IEEE, 2008, pp. 137–142. doi: 10.1109/CDC.2008.4739312.
- [55] F. Itakura and S. Saito, ‘Analysis synthesis telephony based on the maximum likelihood,’ in *Proc. 6th Int. Cong. on Acoust.*, Tokyo, Japan, 1968, pp. C17–C20.
- [56] F. Itakura and S. Saito, ‘A statistical method for estimation of speech spectral density and formant frequencies,’ *Electron. Commun. Jpn.*, vol. 53-A, pp. 36–43, 1970.
- [57] B. Iser, W. Minker and G. Schmidt, Eds., *Bandwidth Extension of Speech Signals* (Lecture Notes in Electrical Engineering). New York, NY, USA: Springer, 2008, ISBN: 987-0-387-68899-2. doi: 10.1007/978-0-387-68899-2.
- [58] P. Jain, B. J. Kulis and I. S. Dhillon, ‘Online Linear Regression using Burg Entropy,’ The University of Texas at Austin, Austin, TX, USA, Tech. Rep. TR-07-08, 2007.

- [59] R. Gray, A. Buzo, A. Gray and Y. Matsuyama, ‘Distortion measures for speech processing,’ *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 28, no. 4, pp. 367–376, 1980. DOI: 10.1109/TASSP.1980.1163421.
- [60] P. T. FitzRoy, ‘Review: [Untitled],’ *J. Mark. Res.*, vol. 7, no. 3, pp. 401–403, 1970, ISSN: 00222437. DOI: 10.2307/3150305.
- [61] X. Jiang, L. Ning and T. T. Georgiou, ‘Distances and Riemannian Metrics for Multivariate Spectral Densities,’ *IEEE Trans. Autom. Control*, vol. 57, no. 7, pp. 1723–1735, 2012. DOI: 10.1109/TAC.2012.2183171.
- [62] D. Chen, M. Davies, M. J. Ehrhardt, C.-B. Schönlieb, F. Sherry and J. Tachella, ‘Imaging With Equivariant Deep Learning: From unrolled network design to fully unsupervised learning,’ *IEEE Signal Process. Mag.*, vol. 40, no. 1, pp. 134–147, 2023. DOI: 10.1109/MSP.2022.3205430.
- [63] K. J. Cheoi, H. Choi and J. Ko, ‘Empirical Remarks on the Translational Equivariance of Convolutional Layers,’ *Appl. Sci.*, vol. 10, no. 9, 3161, 2020. DOI: 10.3390/app10093161.
- [64] O. Ronneberger, P. Fischer and T. Brox, ‘U-Net: Convolutional Networks for Biomedical Image Segmentation,’ in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells and A. F. Frangi, Eds., Cham, Switzerland: Springer International Publishing, 2015, pp. 234–241, ISBN: 978-3-319-24574-4. DOI: 10.1007/978-3-319-24574-4_28.
- [65] G. Xu, X. Wang, X. Wu, X. Leng and Y. Xu, ‘Development of residual learning in deep neural networks for computer vision: A survey,’ *Eng. Appl. Artif. Intell.*, vol. 142, p. 109890, 2025, ISSN: 0952-1976. DOI: <https://doi.org/10.1016/j.engappai.2024.109890>.
- [66] PyTorch Contributors, *Conv3d – PyTorch 2.7 documentation*, <https://docs.pytorch.org/docs/stable/generated/torch.nn.Conv3d.html>, Accessed: 2025-05-13, n.d.
- [67] S. Ioffe and C. Szegedy, ‘Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,’ in *Proc. 32nd Int. Conf. on Mach. Learn.*, F. Bach and D. Blei, Eds., ser. PMLR, vol. 37, Lille, France: PMLR, 2015, pp. 448–456. [Online]. Available: <https://proceedings.mlr.press/v37/ioffe15.html>.
- [68] C. Trabelsi, O. Bilaniuk, Y. Zhang *et al.*, ‘Deep Complex Networks,’ in *Int. Conf. Learn. Represent.*, 2018. [Online]. Available: <https://openreview.net/forum?id=H1T2hmZAb>.
- [69] J. A. Barrachina, C. Ren, G. Vieillard, C. Morisseau and J.-P. Ovarlez, *Theory and Implementation of Complex-Valued Neural Networks*, 2023. DOI: 10.48550/arXiv.2302.08286. arXiv: 2302.08286 [stat.ML].
- [70] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, ‘Dropout: A Simple Way to Prevent Neural Networks from Overfitting,’ *J. of Mach. Learn. Res.*, vol. 15, no. 56, pp. 1929–1958, 2014. [Online]. Available: <http://jmlr.org/papers/v15/srivastava14a.html>.
- [71] S. Park and N. Kwak, ‘Analysis on the Dropout Effect in Convolutional Neural Networks,’ in *Computer Vision – ACCV 2016*, S.-H. Lai, V. Lepetit, K. Nishino

Bibliography

- and Y. Sato, Eds., Cham, Switzerland: Springer International Publishing, 2017, pp. 189–204, ISBN: 978-3-319-54184-6. DOI: 10.1007/978-3-319-54184-6_12.
- [72] M. D. Zeiler, D. Krishnan, G. W. Taylor and R. Fergus, ‘Deconvolutional networks,’ in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, San Francisco, CA, USA: IEEE, 2010, pp. 2528–2535. DOI: 10.1109/CVPR.2010.5539957.
- [73] ‘3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Spatial channel model for Multiple Input Multiple Output (MIMO) simulations (Release 18),’ 3GPP, Valbonne, France, Tech. Rep. 25.996, version 18.0.0, 2024. [Online]. Available: <http://www.3gpp.org>.
- [74] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 2017. DOI: 10.48550/arXiv.1412.6980. arXiv: 1412.6980 [cs.LG].
- [75] S. Bock and M. Weiß, ‘A Proof of Local Convergence for the Adam Optimizer,’ in *Int. Jt. Conf. Neural Netw.*, Budapest, Hungary: IEEE, 2019, pp. 1–8. DOI: 10.1109/IJCNN.2019.8852239.

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden
www.chalmers.se



CHALMERS
UNIVERSITY OF TECHNOLOGY