



Vrije Universiteit Brussel
Applied Sciences and Engineering
Brussels, Belgium



Royal Military Academy
Polytechnic Faculty
Brussels, Belgium

COMPRESSED SENSING FOR MICROWAVE IN-DEPTH IMAGING

Mathias Becquaert

Supervisors:
Prof Dr Ir Marijke Vandewal,
Prof Dr Ir Johan Stiens

Thesis submitted in fulfillment of the requirements for the degree of Doctor in
Engineering Sciences (RMA) and Doctor in Engineering Sciences (VUB)

May 3th 2019

Alle Rechten voorbehouden. Niets van deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook, zonder voorafgaande schriftelijke toestemming van de auteur.

All rights reserved. No part of this publication may be produced in any form by print, photoprint, microfilm, electronic or any other means without permission from the author.

Printed by
Crazy Copy Center Productions
VUB Pleinlaan 2, 1050 Brussel
Tel / fax : +32 2 629 33 44
crazycopy@vub.ac.be
www.crazycopy.be

ISBN : 9789493079229

NUR CODE : 950

Abstract

Low hardware costs, short measurement times, fast data handling, and flexibility in taking measurements under harsh conditions, are some of the key requirements of today's sensors. These sensor characteristics are influenced or determined by the number of measurements. The classic sampling theory fixes a lower bound on the number of samples needed, known as the Shannon-Nyquist sampling rate. Over the past decade, a new measurement methodology has gained a lot of attention: Compressed Sensing (CS). The advent of CS is motivated by the observation that in many applications the measurements, performed at high sampling rates and thus generating high data volumes, are followed by a compression step. Compressed Sensing incorporates the compression step into the measurement phase. Under certain conditions on the signal sparsity and on the measurement methodology, CS allows to reconstruct the sensed signal from a number of samples far below the Shannon-Nyquist bound. In this work, Compressed Sensing is applied on Stepped-Frequency Continuous Wave (SFCW) and Synthetic Aperture Radar (SAR) measurements. In our work, the applicability of Compressed Sensing on SFCW and SAR data is evaluated for two applications: (1) The Non-Destructive Testing (NDT) of 3D-printed objects and (2) Through-the-Wall Radar Imaging (TWRI).

3D-printing or Additive Manufacturing (AM) is a rapid growing production technique which allows to manufacture any part from a Computer-Aided Design (CAD) model. Defects and printing flaws are common in AM and the quality of the printed objects can vary. An evaluation of the integrity of the inner structure of the part is thus necessary. In our study, we explore a NDT technique based on a mm-wave sensor using a SAR architecture for polymer printed objects. The proposed NDT technique has the advantage, compared to common techniques like X-ray tomography or ultrasound, to be contactless and to use non-ionizing radiation. The technique performs a lot of time-consuming measurements, generating large amounts of data. Compressed Sensing is utilized in this application to decrease the number of samples. The number of samples is further decreased by adding Side Information (SI) which can be a previous measurement of the same or a similar object or a synthetic measurement.

Through-the-Wall Radar Imaging (TWRI) is gaining a lot of interest by both the research community and the industry and is of interest for police, defence forces, fire and rescue personnel and first responders. A TWRI radar utilizes UWB signals combined with a scanning strategy along a wall to image the hidden scene and detect the presence of human targets behind the illuminated wall. A large technological gap exists between the commercial solutions, which are fast and easy to handle, and the published research prototypes, which deliver high resolution images. Compressed Sensing seems a good candidate to reduce that gap. In the presented study, a promising approach is tested. We combine different CS techniques which results in a sensor capable to deliver high quality TWRI radar images from a low number of samples by intelligently exploiting the already scanned part of the scene. Moreover, the sensor autonomously adapts his sampling rate to the signal sensed along the scanning path. This new approach is successfully tested on real TWRI data for the detection of static human targets through a wall.

Acknowledgements

Firstly, I would like to thank my PhD supervisors: Prof Dr Ir Marijke Vandewal and Prof Dr Ir Johan Stiens. Prof Vandewal has given me the opportunity to start with this adventure. She has been supportive all along this journey by giving advice, sharing ideas, reviewing my work and so much more. Thank you for everything. I have learned so much from you in many domains. Prof Stiens, thank you for helping me realising this work. I especially would like to thank you for helping me with the experiments at the ETRO Department.

My profound gratitude goes to Prof Dr Ir Nikos Deligiannis. Thank you for sharing with me your visions, opinions and enthusiasm during the many meetings we had together.

I would also thank my colleagues of the CISS department of the Royal Military Academy for being supportive during this work. I would especially like to thank Dr Ir Edison Cristofani for the many insightful discussions we had on the topic of Compressed Sensing and for all the papers we published together.

Last but not least I want to thank my friends and family. My profound gratitude goes to my wife Barbara Roobroeck, who stands by my side without hesitation. Thank you for your love, for being comprehensive and for your enthusiasm. Each time I have encountered a difficulty during this adventure you came up with a solution and helped me forward. Thank you for everything!

Members of the Jury

Prof Dr Ir PERNEEL, Christiaan (President)
Royal Military Academy, MWMW Department, BE

Prof Dr Ir PINTELON, Rik (Vice president)
Vrije Universiteit Brussel, ELEC Department, BE

Prof Dr Ir HAELTERMAN, ROBBY (Secretary)
Royal Military Academy, MWMW Department, BE

Prof Dr Ir VANDEWAL, Marijke (Promoter)
Royal Military Academy, CISS Department, BE

Prof Dr Ir STIENS, Johan (Promoter)
Vrije Universiteit Brussel, ETRO Department, BE

Prof Dr MARTORELLA, Marco
University of Pisa, Department of Information Engineering, IT

Prof Dr Ir PIZURICA, Aleksandra
Universiteit Gent, TELIN Department, BE

Prof Dr Ir DELIGIANNIS, Nikolaos
Vrije Universiteit Brussel, ETRO Department, BE

Dr Ing. WEISS, Matthias
Fraunhofer FHR, D

Contents

1	Introduction	1
	Research questions	2
	Unique contributions of this thesis	3
	Organization of this dissertation	3
2	Compressed Sensing	5
2.1	Introduction	5
2.2	The basic problem	5
2.2.1	A set of linear measurements	5
2.2.2	The Shannon-Nyquist theorem	6
2.2.3	Undersampling	6
2.3	Main results in Compressed Sensing	7
2.3.1	Conditions on the signal	7
2.3.2	Conditions on the measurement matrix	8
	Null-space property	8
	Coherence property	8
	Restricted Isometry Property	9
2.3.3	CS reconstruction algorithms	10
	Basis Pursuit algorithms	10
	Matching Pursuit algorithms	11
	Matching Pursuit versus Basis Pursuit algorithms	13
3	Compressed Sensing applied on SFCW and SAR measurements	19
3.1	Introduction	19
3.2	SFCW measurements	20
3.2.1	The Discrete Fourier Transform	20
3.2.2	Stepped-Frequency Continuous Wave measurements	22
3.2.3	CS applied on SFCW measurements	24
	CS reconstruction evaluation	25
	Sampling strategies	28
3.3	SAR measurements	29
	The SAR measurement matrix	31
	Phase diagram for SAR measurements.	32
3.4	Sequential Compressed Sensing	33
3.4.1	Derivation of an upper bound for the reconstruction error	34
3.4.2	Influence of the approximations	36
3.4.3	An example	41
3.4.4	Empiric evaluation of the proposed method	42

3.4.5	Number of cross-validation measurements and the upper bound probability	44
3.5	Preconditioning methods	48
3.5.1	Singular Value Decomposition	48
3.5.2	Principal Component Analysis	49
3.6	Compressed Sensing with Side Information	50
3.6.1	Homogeneous Side Information	50
3.6.2	Heterogeneous Side Information	55
	Coupled dictionary learning	55
	CS with heterogeneous SI	57
	CS with heterogeneous SI for change detection	60
3.7	Conclusions	61
4	Non-Destructive Testing of polymer 3D-printed objects	63
4.1	Introduction	63
4.2	NDT experiment	68
4.3	CS reconstruction	72
4.3.1	CS reconstruction versus traditional SAR reconstruction	72
4.3.2	CS with side information	76
	CS with Coherent background subtraction	76
	$\ell_1\ell_1$ -minimization	78
	CS with Weighted side information	79
4.3.3	CS with multiple homogeneous weighted SI	85
4.4	Conclusions	88
5	Through-the-Wall Imaging	89
5.1	Introduction	89
5.1.1	State of the art research on Through-the-Wall Imaging	89
	Wall modelling	90
	Wall mitigation techniques	90
	Multipath	91
	Compressed Sensing	91
5.1.2	Commercial Through-the-Wall Imaging sensors	92
5.2	Through-the-Wall radar experiments	97
5.2.1	Through-the-Wall radar prototype	97
	Description	97
	TWI radar parameters	99
5.2.2	Measurement setup	101
5.2.3	Classic Synthetic Aperture Radar	103
5.3	Compressed Sensing reconstruction	104
5.3.1	Through-the-Wall radar measurement matrix	104
	Wall ringing multipath and propagation delay	105
	Wall transmission coefficients	106
	Antenna pattern	106
	Received signal model	109
5.3.2	Compressed sensing reconstruction	109
	Evaluation of the reconstruction performance	111
5.3.3	Sampling strategies	113

5.3.4	Sparsifying bases	115
5.3.5	Wall mitigation techniques	119
	(1) Coherent background subtraction:	120
	(2) Moving average filter:	122
	(3) Singular value decomposition:	124
	(4) Robust Principal Component Analysis:	125
5.4	Range profile reconstruction	127
5.4.1	Single range profile reconstruction	128
5.4.2	Side information	131
5.4.3	Multiple weighted side information	132
5.5	Sequential Compressed Sensing	137
5.6	Conclusions	141
6	Conclusions	143
6.1	Conclusions of the different chapters	143
6.2	Final conclusion	145
6.3	Future work	145

Publications

Articles in scientific journals:

- Becquaert, M., Cristofani, E., Lauwens, B., Vandewal, M., Stiens, J., Deligiannis, N. (2019). Online Sequential Compressed Sensing with Multiple Information for Through-the-Wall Radar Imaging. *IEEE Sensors*, Accepted: January 2019.
- Becquaert, M., Cristofani, E., Van Luong, H., Vandewal, M., Stiens, J., Deligiannis, N. (2018). Compressed Sensing mm-Wave SAR for Non-Destructive Testing Applications Using Multiple Weighted Side Information. *Sensors (Basel, Switzerland)*, 18(6).
- Cristofani, E., Becquaert, M., Lambot, S., Vandewal, M., Stiens, J., Deligiannis, N. (2018). Random Subsampling and Data Preconditioning for Ground Penetrating Radars. *IEEE Access*.
- Pourkazemi, A., Stiens, J. H., Becquaert, M., Vandewal, M. (2017). Transient Radar Method: Novel Illumination and Blind Electromagnetic/Geometrical Parameter Extraction Technique for Multilayer Structures. *IEEE Transactions on Microwave Theory and Techniques*, 65(6), 2171-2184.
- Brook, A., Cristofani, E., Becquaert, M., Lauwens, B., Vandewal, M. (2013). Applicability of compressive sensing on three-dimensional terahertz imagery for in-depth object defect detection and recognition using a dedicated semisupervised image processing methodology. *Journal of Electronic Imaging*, 22(2), 021004.
- Cristofani, E., Becquaert, M., Vandewal, M. (2013). Performance of 2D compressive sensing on wide-beam through-the-wall imaging. *Journal of Electrical and Computer Engineering*, 2013, 6.

Conference papers:

- Becquaert, M., Cristofani, E., Vandewal, M., Stiens, J., Deligiannis, N. (2018). Online Sequential Compressed Sensing with Weighted Multiple Side Information for Through the Wall Imaging. In 27th URSI Benelux Forum, Delft, the Netherlands, 2018. URSI.
- Becquaert, M., Cristofani, E., Stiens, J., Vandewal, M., Deligiannis, N. (2017). Compressed Sensing SAR Through-the-Wall Imaging with Side Information. In 26th URSI Benelux Forum, Brussels, Belgium, 2017. URSI.
- Becquaert, M., Cristofani, E., Pandey, G., Vandewal, M., Stiens, J., Deligiannis, N. (2016, May). Compressed sensing mm-wave SAR for non-destructive testing applications using side information. In Radar Conference (RadarConf), 2016 IEEE (pp. 1-5). IEEE.

- Becquaert, M., Cristofani, E., Vandewal, M. (2013, October). On the applicability of compressive sensing on FMCW synthetic aperture radar data for sparse scene recovery. In Radar Conference (EuRAD), 2013 European (pp. 9-12). IEEE.
- Cristofani, E., Mahfoudia, O., Becquaert, M., Neyt, X., Horlin, F., Deligiannis, N., Stiens, J., Vandewal, M.(2018). Exploring Side Information for DVB-t-based Passive Radars. In 27th URSI Benelux Forum, Delft, the Netherlands, 2018. URSI.
- Cristofani, E., Mahfoudia, O., Becquaert, M., Neyt, X., Horlin, F., Deligiannis, N., Stiens, J., Vandewal, M.(2017). Compressive Sensing and DVB-T-Based Passive Coherent Location. In 26th URSI Benelux Forum, Brussels, Belgium, 2017. URSI.
- Cristofani, E., Becquaert, M., Pandey, G., Vandewal, M., Deligiannis, N., Stiens, J. (2016, September). Compressed sensing and defect-based dictionaries for characteristics extraction in mm-Wave non-destructive testing. In Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2016 41st International Conference on (pp. 1-2). IEEE.

List of Figures

2.1	Mean reconstruction time (over 64 iterations) for increasing signal lengths (N) for random vectors ($k/N = 5\%$ and $n/N = 30\%$). The reconstructions are obtained with (1) the BP algorithms: SPGL1, ℓ_1 -magic and CVX and (2) the MP algorithms: OMP and COSAMP.	14
2.2	Mean reconstruction time (over 64 iterations) for increasing sparsity rates (k/N) for random vectors ($N = 512$ and $n/N = 30\%$). The reconstructions are obtained with (1) the BP algorithms: SPGL1, ℓ_1 -magic and CVX and (2) the MP algorithms: OMP and COSAMP.	14
2.3	Boxplots of the relative error for BP (CVX) and MP (OMP and CoSaMP) algorithms. In (a) : reconstruction from noiseless measurements and in (b): reconstruction from measurements with a SNR = 15 dB.	16
2.4	Phase diagrams using a BP algorithm (SPGL1), OMP and CoSaMP. On the left hand side: the mean relative reconstruction error as a function of the undersampling fraction δ and the sparsity ratio ρ . On the right hand side: the computation time as a function of δ and ρ	17
3.1	The upper graph: the original signal (in blue) with in red: the selected samples at a sampling rate of 2 kHz and in yellow the samples selected at a sampling rate equal to 1 kHz. The middle graph: the spectrum of the signal sampled at 2 kHz. The lower graph: the spectrum of the signal sampled at 1 kHz.	21
3.2	The upper graph: the original signal (in blue) with marked in red: the equidistant samples and in yellow: the randomly selected samples. The middle graph: the spectrum of the reconstructed signal from the equidistant samples. The lower graph: the spectrum of the reconstructed signal from the randomly selected samples.	22
3.3	Stepped-frequency continuous wave signal.	23
3.4	Reconstruction error obtained after the reconstruction of a signal from a sampled SFCW measurement, for an increasing number of samples, by applying an IFFT (in blue) or a CS reconstruction algorithm (SPGL1) (in red).	25
3.5	Phase diagram obtained using SFCW measurements with SPGL-1. The color scheme depicts the probability of a successful reconstruction. We observe two distinct areas: the blue area where the probability is equal to 0 and the red area where the probability is equal to 1.	26

3.6	Phase diagram obtained using SFCW measurements with CVX. The color scheme depicts the probability of a successful reconstruction. (blue: probability = 0; red: probability = 1).	27
3.7	Phase diagram obtained using SFCW measurements with OMP. The color scheme depicts the probability of a successful reconstruction. (blue: probability = 0; red: probability = 1).	28
3.8	Reconstruction success rate (with the confidence intervals) for different sampling strategies ($N = 840$, and $n = 29$).	29
3.9	Schematic representation of the working principal of a SAR with physical opening angle θ and synthetic aperture L	30
3.10	Schematic representation of the scene sensed to fill the columns of A corresponding to the pixel x_{ij}	31
3.11	Phase diagram for SAR measurements with SPGL1. The color scheme depicts the mean relative reconstruction error (blue: reconstruction error is very small, red: high reconstruction error).	32
3.12	Phase diagrams depicting the outcome of the KS-test on the null-hypothesis that the expressions in (3.32) and (3.33) are drawn from a χ^2_{2T} distribution. In (a) and (b) are respectively the test decisions (1: rejects the null-hypothesis and 0: accept the null-hypothesis) and the mean p -values (over 30 executions) for (3.32). In (c) and (d) are the test decisions and the p -values for (3.33).	38
3.13	In (a) the PDFs and in (b) the corresponding CDFs, with in blue: the PDF/CDF of the theoretical χ^2_{2T} distribution with $T = 4$. In red: the PDF/CDF obtained from the expression in (3.32). In yellow: the PDF/CDF obtained when replacing $\hat{\mathbf{A}}$ with \mathbf{A} in (3.32). In magenta: the PDF/CDF obtained when replacing $\hat{\mathbf{A}}$ with \mathbf{A} in (3.32) and with the approximations on $E[\mathbf{z}_i]$ and $\text{var}(\mathbf{z}_i)$	40
3.14	Reconstruction error for increasing sampling rates (in blue: $k = 60$ and in red: $k = 120$). The corresponding upper bounds are represented by the dashed curves.	41
3.15	(a) Phase diagram with the true reconstruction error. (b) Estimated upper bound for the reconstruction error (using $T = 10$ cross-validation measurements and $\alpha = 0.05$). (c) Difference between the true reconstruction error and the estimated upper bounds. (d) In red: combinations of ρ and δ where the upper bound is lower than the true error.	43
3.16	(a) Phase diagram with the mean reconstruction error for all combinations of ρ and δ over 30 iterations. (b) Phase diagram with the mean of the upper bounds over the 30 iterations. (c) Phase diagram with the mean of the difference between the true error and the upper bound. (d) Phase diagram with the variance of the difference between the true error and the upper bound.	44
3.17	(a) The measured true error δ (solid blue curve) and the estimated bounds for an increasing number of measurements using a different number of cross validation measurements (T). (b) Zoom of figure (a)	45

3.18	(a) CDF's of the χ_T^2 distribution function for increasing T and the z^* for $\alpha = 0.1$ are encircled in red. (b) True measured error, the estimated error and the estimated error bounds obtained with the z^* from (a) with $\alpha = 0.1, 0.01, 0.05$ and 0.2	47
3.19	Mean reconstruction error and confidence intervals for increasing subsampling rates obtained with CS, CS with background subtraction and RAMSIA. (a) $k_x \leq k_z$ and (b) $k_x \geq k_z$	52
3.20	Mean reconstruction error for increasing subsampling rates obtained with CS (the blue curve), CS with coherent background subtraction (the red curve) and RAMSIA (the yellow curve) using a single SI with different quality levels (change rates: 10%, 15%, 20%, 25%, 30% and 99%).	53
3.21	Mean reconstruction error and confidence intervals for increasing subsampling rates obtained with (a) CS (the blue curve) and CS with background subtraction and (b) CS (the blue curve) and RAMSIA for multiple SIs with varying change rates. (SI1) $J = 1$ (change rate 5%); (SI2) $J = 2$ (change rates = 5% and 5%); (SI3) $J = 2$ (change rates = 5% and 20%).	54
3.22	Dictionary identifiability of the coupled dictionaries Ψ^c , Φ^c and Φ learned from an increasing number of training measurements.	57
3.23	Reconstruction error for increasing subsampling rates when reconstructing the original signal with CS without SI (dashed line) and with hetero-SI (full line) for different values of s_v	58
3.24	Reconstruction error for increasing subsampling rates when reconstructing the original signal with CS without SI (dashed line) and with hetero-SI (full line) for different values of s_z	59
3.25	Reconstruction error for increasing subsampling rates when reconstructing the original signal with CS without SI (dashed line) and with hetero-SI (full line) using the learned dictionaries from different numbers of training measurements (T).	60
3.26	Reconstruction error for increasing subsampling rates when reconstructing the difference between two heterogeneous measurements using CS with heterogeneous SI.	61
4.1	Examples of defects in 3D printed parts obtained with a single FDM machine.	65
4.2	Scheme of the measurement setup for the ISAR NDT experiments.	69
4.3	(a) picture taken during a NDT measurement of the test object. (b) schematic representation of the test object with the three defects.	70
4.4	mm-wave ISAR images (horizontal cut) of the test object in (a) without defects, in (b) with 1 defect, in (c) with 2 defects and in (d) with 3 defects.	71
4.5	Histograms of the normalized pixel values. (a) Histogram of the ISAR image to reconstruct. (b) Histogram of the DCT of the ISAR image to reconstruct.	74
4.6	reconstruction for different subsampling rates. Left column: images obtained with traditional SAR algorithm, middle column: results obtained with CS and right column: results obtained with CS and the DCT transform.	75

4.7	Mean reconstruction error and confidence intervals for increasing subsampling rates obtained with a standard SAR algorithm, with a CS algorithm (SPGL1) and with a CS algorithm with a DCT transform. The dashed curves are the corresponding standard deviations.	76
4.8	Mean reconstruction error for increasing subsampling rates obtained from the CS SAR image reconstruction with coherent background subtraction without SI, with SI 1, SI 2 and SI 3.	77
4.9	Mean reconstruction error for increasing subsampling rates obtained from the CS SAR image reconstruction with $\ell_1\ell_1$ -minimization without SI, with SI 1, SI 2 and SI 3.	79
4.10	inter SI weights β_1 using a single SI (SI 1, SI 2 or SI 3), for increasing subsampling rates.	80
4.11	(a) Absolute value of the difference between the SAR image to reconstruct and the SI 1 over all the elements of the SAR image. (b) The intra SI weights obtained with RAMSIA for a subsampling rate equal to 90%. (c)The intra SI weights (SI = SI 1) obtained with RAMSIA for a subsampling rate equal to 30%.	82
4.12	(a) Absolute value of the difference between the SAR image to reconstruct and the SI 3 over all the elements of the SAR image. (b) The intra SI weights obtained with RAMSIA for a subsampling rate equal to 90%. (c)The intra SI weights (SI = SI 3) obtained with RAMSIA for a subsampling rate equal to 30%.	83
4.13	Mean reconstruction error for increasing subsampling rates obtained with the CS SAR image reconstruction with weighted SI: without SI, with SI 1, SI 2 and SI 3.	84
4.14	Reconstructed SAR images obtained with coherent background subtraction (left column), $\ell_1\ell_1$ -minimization (middle column) and RAMSIA (right column) for subsampling rates equal to 60%, 40%, 30% and 20%.	86
4.15	Mean reconstruction error for increasing subsampling rates obtained with coherent background subtraction, $\ell_1\ell_1$ -minimization and RAMSIA using the three SIs	87
4.16	inter SI weights β_j for increasing subsampling rates obtained with the RAMSIA algorithm.	87
5.1	(a) A picture of the through-the-wall system setup with (1) the motorized board, (2) the VNA, (3) the power source, (4) the horn antenna and (5) the computer. (b) A screenshot of one of the windows of the user interface of the Labview program.	97
5.2	Transmission coefficient S_{21} in [dB] for free space, with an aerated concrete wall and with a hollow concrete wall between the two antennas.	100
5.3	(a) Aerated concrete block and (b) Hollow concrete block used to build the walls for the TWI experiments.	101
5.4	Example of an experiment with the TWI radar system.	102

5.5	Images obtained with a Range Doppler SAR compression algorithm of three scenes together with a schematic representation of the scenes.	103
5.6	Round trip paths with maximum 1 reverberation and 1 multipath reflection.	105
5.7	Schematic representation of the EM signal path through a wall.	105
5.8	Antenna diagrams at 1 GHz to 5 GHz.	108
5.9	CS reconstruction (using l1-magic) for undersampling rates ranging from 100% down to 50%.	110
5.10	SAR image of the scene with framed in red: the target areas and in blue: the clutter areas.	111
5.11	PSNR of the CS reconstructed TWI measurement as a function of the subsampling ratio	112
5.12	Sampling strategies	114
5.13	PSNR of the CS reconstructed TWI obtained with four different sampling strategies.	115
5.14	Histograms of the normalized pixel values of (a) the original TWI scene, (b) the wavelet transform of the scene and (c) the DCT of the scene.	116
5.16	CS reconstruction, left column: without transform, middle column: using a wavelet transform and right column: using the DCT.	118
5.17	PSNR for increasing subsampling rates without using a sparsifying transform, with the wavelet transform and with the DCT.	118
5.18	Reconstruction of the TWI scene with a single target for subsampling rates equal to 100%, 80%, 70% and 60%.	120
5.19	CS Reconstruction with coherent background subtraction for sampling rates equal to 100%, 50%, 25% and 10%.	121
5.20	Reconstruction of the TWI scene with a single target, after applying a moving average filter on the raw data, for subsampling rates equal to 100%, 60%, 50% and 40%.	122
5.21	Reconstruction of the TWI scene with a single target using SVD for subsampling rates equal to 100%, 80%, 70% and 60%.	124
5.22	RPCA reconstruction of the TWI measurement with on the left: the background and on the right: the foreground with the target.	125
5.23	Reconstruction of the small scene with a single target using RPCA obtained using 100%, 40%, 30% and 20% of the original number of samples.	126
5.24	flowcharts of TWI CS approaches. On the left: CS SAR image reconstruction using a SAR measurement matrix. On the right: TWI obtained after a CS reconstruction of range profiles, using a DFT measurement matrix, and cross-range compression.	127
5.25	(a) Range compressed TWI with range range profile nr 42 framed with the red rectangle. (b) Plot of the range profile at position 42.	129
5.26	(a) CS reconstruction of range profile 42 using 100%, 90%, 80% and 70% of the original raw data. (b) Mean reconstruction error for increasing subsampling rates.	130

5.27	PSNR for the CS SAR image reconstruction (in blue) and the TWI obtained from the CS range profile reconstruction (in red) for increasing subsampling rates.	131
5.28	(a) Range profiles at sensor positions 1 (red curve) and 42 (blue curve). (b) Mean reconstruction error for the CS reconstruction of range profile 42 without SI (blue curve) and with SI (red curve).	132
5.29	(a) Resulting range profile after subtracting the range profile at position 1 from the range profile at position 42. (b) Corresponding intra SI weights attributed by the RAMSIA algorithm.	133
5.30	(a) Euclidean distance between the range profile at position 78 and all the previous range profiles. (b) attributed inter SI weights to the range profiles for the reconstruction of range profile 78.	134
5.31	Resulting images after the reconstruction of the range profiles with RAMSIA for subsampling rates equal to 50%, 40%, 30% and 20%.	135
5.32	Reconstruction error [dB] of the range profiles over all sensor positions (vertical axis) for increasing subsampling rates (vertical axis). (a) without SI (b) ℓ_1 -reconstruction and (c) RAMSIA.	136
5.33	Flowchart explaining the approach combining multiple weighted side information and sequential compressed sensing.	137
5.34	True error (in blue) and estimated error (in red) for in (a) using sequential CS without SI and in (b) using sequential CS with weighted SI.	139
5.35	(a) True error [dB] and (b) the estimated error bounds [dB] for the reconstruction of all the range profiles along the scanning path for increasing subsampling rates.	140
5.36	(a) the reconstructed range profiles and (b) the reconstructed TWI SAR image by combining RAMSIA and sequential CS.	141

List of Tables

4.1	68
4.2	81
5.1	Overview of TWI radar prototypes used in a research context	93
5.2	Overview of commercial of-the-shelf TWI radars	95
5.3	Overview of commercial of-the-shelf TWI radars - continued	96
5.4	99
5.5	TCR for the three targets.	112

List of abbreviations

ABS	A crylonitile B utadiene S tylene
AM	A dditive M anufacturing
BAAM	B ig A rea A dditive M anufacturing
BP	B asis P ursuit
BPDN	B asis P ursuit D e N oising
BS	B ackground S ubtraction
CA-CFAR	C ell A veriging - C onstant F alse A larm R ate
CAD	C omputer A ided D esign
CDF	C umulative D istribution F unction
CMP	C omplementary M atching P ursuit
CoSaMP	C ompressive S ampling M atching P ursuit
COTS	C ommercial O f- T he- S helf
CS	C ompressed S ensing
DCP	D isciplined C onvex P rogramming
DCT	D iscrete C osine T ransform
DFT	D iscrete F ourier T ransform
DWT	D iscrete W avelet T ransform
EM	E lectro M agnetic
FDM	F used D eposition M odelling
FIR	F inite I mpulse R esponse
FISTA	F ast I terative S hrinkage- T hresholding A lgorithm
FMCW	F requency M odulated C ontinuous W ave
GPR	G round P enetrating R adar
IFFT	I nverse F ast F ourier T ransform
ISAR	I nverse S ynthetic A perture R adar
LASSO	L east A bsolute S hrinkage and S election O perator
MA	M oving A verage
MOMP	M odified O rthogonal M atching P ursuit
MRI	M agnetic R esonance I magery
MVNA	M illimeter V ector N etwork A nalyser
NDE	N on- D estructive E valuation
NDT	N on- D estructive T esting
NP	N on-deterministic P olynomial-time
OMP	O rthogonal M atching P ursuit

PCA	Principal Component Analysis
PDF	Probability Density Function
PSNR	Peak Signal-to-Noise Ratio
RAMSIA	Reconstruction Algorithm with Multiple SI using Adaptive Weights
RIP	Restricted Isometry Property
RPCA	Robust Principal Component Analysis
RTT	Round Trip Time
SAR	Synthetic Aperture Radar
SFCW	Stepped Frequency Continuous Wave
SI	Side Information
SLS	Selective Laser Sintering
SNR	Signal-to-Noise Ratio
StOMP	Stagewise Orthogonal Matching Pursuit
SVD	Singular Value Decomposition
TCR	Target-to-Clutter Ratio
TEM	Transverse Electromagnetic Mode
TW(R)I	Through-the-Wall (Radar) Imaging
VNA	Vector Network Analyser

Chapter 1

Introduction

The number of sensing devices is exponentially growing as well in almost all branches of the industry as in our daily live. A good example is the massive introduction of sensors in the automotive industry. Today's modern cars have an average of 60 up to 100 sensors on board. This steep growth in number of sensors and applications is only made possible thanks to the ongoing digital revolution. The digitization allows the design and production of cheaper, more flexible and more robust sensing devices when compared to their analog counterparts [Davenport et al.2011]. At the very heart of this revolution lies the theoretical work of Kotelnikov, Nyquist, Shannon and Whittaker [Kotelnikov1933, Nyquist1928, Shannon1949, Whittaker1915]. Their results dictate that the information of an analog signal can be fully captured by taking equidistant samples at a sampling rate equal to at least two times the highest frequency present within the analog signal. Since then, this theorem, commonly referred to as the Shannon-Nyquist theorem, has governed the design of all digital sensing systems. But, this theoretical bound entails practical and technical challenges.

First, in some applications, as for example in imaging and video applications, measuring at the imposed Nyquist rate, generates vast volumes of data that are difficult and cumbersome to handle and to transmit. This problem has been tackled with the rise of the transform coding schemes which form the foundations of for example the JPEG2000 [Acharya and Tsai2005], the MPEG [Le Gall1991] or the MP3 [McCandless1999] standards. The aim of these lossy compression techniques is to represent the high-dimensional data in a concise way, while limiting the distortions introduced by doing so. Therefore, a transformation basis, such as the Discrete Cosine Transform (DCT) (used in the JPEG standard) or the wavelet transform (used in the JPEG2000 standard), is needed. These transformations allow to concentrate the important information of the signal in a small number of coefficients which are preserved, whereas the rest of the coefficients are neglected.

A second challenge is that in some applications sampling at the Nyquist rate is technically impossible (because of, for example, the limitations of analog-to-digital converters [Walden1999]) or very costly or time consuming. This problem is not addressed by the compression techniques. A full measurement is first performed at a sampling rate dictated by the Nyquist rate, before the compression is applied and the data is reduced by throwing away that part that was measured but appears not to be relevant. The idea of incorporating the compression phase during the acquisition gained a lot of attention with the work of Candès, Romberg, Tao and Donoho [Candès et al.2006, Donoho2006]

around 2006. This new framework is known as Compressed Sensing (CS) or Compressive Sampling. Their work started an avalanche of publications in a multitude of research fields and applications. One of the first impressive results were obtained in the field of medical imaging. More precisely, the application of CS on Magnetic Resonance Imagery (MRI) can reduce the measurement time to a fraction of the original required time [Lustig et al.2007, Gamper et al.2008]. Encouraged by these many and relevant results, we started our study on applying CS in the field of microwave imaging.

Research questions

In this dissertation we try to explain and evaluate different approaches for applying CS in the domain of microwave imaging, using synthetic and real measurements. We will search and formulate answers to the following series of subjects and questions.

1. The microwave images obtained from as well the simulated as from the real data are obtained with a Vector Network Analyser. This type of equipment measures the reflection parameter at discrete frequencies. This implies that the signal in one dimension corresponds to a Stepped Frequency Continuous Wave (SFCW) signal. This 1D-signal is expanded to a 2D-measurement in order to obtain an image, by using a Synthetic Aperture Radar (SAR) strategy. To start with, we want to evaluate the applicability and measure the impact of applying CS on SFCW and on SAR measurements.
2. Next, we will explore and test recent and promising techniques and extensions on CS using synthetic experiments. More precisely, we want to evaluate the performance of CS combined with: (1) a sequential measurement approach, (2) a Singular Value Decomposition (SVD), (3) a Principal Component Analysis (PCA), (4) (multiple) homogeneous side information(s) and (5) heterogeneous side information.
3. A first application we selected is the Non-Destructive Testing (NDT) of polymer additive manufactured parts. We first want to evaluate if mm-wave SAR sensing can be a relevant alternative for the existing structural health monitoring techniques, such as X-ray tomography or ultrasound sensing, by forming a mm-wave image of the inner structure of 3D-printed objects. Subsequently, we want to evaluate the performance of CS for this specific application. In many NDT cases, complementary information on the object under test is known prior to the measurement. This can for example be a measurement of the same or a similar part earlier in time. We want to explore, evaluate and compare different techniques for incorporating this side information during the CS construction of the mm-wave image, using real data.
4. The second application is Through-the-Wall Imaging (TWI) using a Synthetic Aperture Radar. This application was selected because we observe a huge gap between the TWI radar prototypes presented in a research context and the commercially available systems. We believe that CS is a good candidate to help to close this gap. The most important differences are: the image quality and the image resolution and on the other hand the sensor user-friendliness and measurement times. We performed a series of experiments, with a self-built prototype, by imaging scenes populated with multiple human targets behind a wall. Using the real data obtained with these experiments, we want to evaluate the applicability of CS on this type of real data. We finally want to combine different CS techniques such as CS with multiple weighted SI and sequential CS.

Unique contributions of this thesis

Multiple novel techniques and approaches are presented in this dissertation. The unique contributions obtained in the framework of this PhD are enlisted hereunder:

- If a measured signal is sparse or has a sparse representation, Compressed Sensing reconstruction techniques allow to reconstruct that signal from a number of samples obtained at a rate below the Nyquist rate. The number of samples needed to obtain a correct reconstruction depends, amongst other things, on the sparsity of the signal. The exact sparsity of the signal is unknown and can only be guessed prior to the measurement. In [Malioutov et al.2010] a method is proposed to fix an upper bound on the reconstruction error for Gaussian measurements, using a small number of extra cross-validation samples. In section 3.4, this approach is enlarged to the Discrete Fourier Transform measurements. To estimate the upper bound several approximations are made. We empirically study the impact of the approximations and show that the upper bound will be overestimated but remains relevant. We further demonstrate that the number of cross validation measurements can be kept fairly small compared to the total number of unknowns.
- In Section 3.6.2 a novel approach is presented to perform change detection using heterogeneous SI. The proposed technique is based on a source separation technique, presented in [Deligiannis et al.2017], and consists of a coupled dictionary learning step end a CS reconstruction step with Side Information.
- In the field of (I)SAR imaging, coherent background subtraction is the commonly applied technique to exploit the Side Information coming from a measurement sharing a high degree of similarity with the signal to reconstruct. In Section 4.3.2 a CS reconstruction technique incorporating multiple weighted Side Informations is proposed in stead of the coherent background subtraction technique. The proposed solution has the advantage over the background subtraction method, that it is not affected by low quality Side Information.
- In Section 5.5, a new way of applying CS in Through-the-Wall Radar Imaging is proposed. This novel approach blends the CS-reconstruction phase with the measurement phase. In a nutshell: the SAR sequentially performs measurements until it detects autonomously that a certain reconstruction quality has been obtained (without having access to the ground truth). The reconstruction is performed using a CS algorithm which exploits the already performed part the SAR measurements as weighted Side Information. This novel approach has the advantage that only the desired reconstruction quality needs to be chosen whereas the sensor adapts the sampling rate on the fly in order to achieve that requirement.

Organization of this dissertation

The remainder of this thesis is organized as follows:

- Chapter 2 gives the reader a concise overview of CS and highlights some important theoretical results. More precisely, we discuss the necessary conditions on the measurement modality and on the signal properties for CS to be applicable. We end the chapter with a discussion of different CS reconstruction algorithms which will be used for obtaining the results presented in the later chapters.

- In chapter 3, we explore a set of CS methods and approaches which will be applied on the real data in chapters 4 and 5 of this dissertation. We start by evaluating the applicability of CS on synthetic data obtained through SFCW and SAR measurements. Further in this chapter, we explore multiple novel techniques combined with CS and SFCW and SAR measurements. Those approaches are: sequential CS, PCA, SVD, CS with homogeneous SI, dictionary learning and CS with heterogeneous SI. This chapter can be seen as a toolbox of techniques from which we will select the appropriate approaches in the next chapters.
- Chapter 4 starts with an outline of the NDT-needs for additive manufactured parts. We then explain the measurement setup and the measurements performed on a 3D-printed test object. Subsequently, we evaluate the impact of using a CS reconstruction algorithm on the robustness against sampling below the Nyquist rate. Finally, we compare different approaches for adding homogeneous SI during the CS construction of the NDT image: coherent background subtraction, $\ell_1\ell_1$ -minimization and CS with multiple weighted SI.
- In chapter 5, we discuss the use of CS on real TWI data. The chapter starts with a comparison between the research TWI radar prototypes and the commercial systems. Next, we build our own TWI radar prototype, based on a Vector Network Analyser. This system is used to perform a series of measurements. These measurements are used to compare several wall mitigation techniques combined with CS. Finally, we propose an online reconstruction technique combining a sequential reconstruction approach with multiple weighted side information.
- Chapter 6 concludes this dissertation. In this final chapter, we resume our results and list the most important contributions obtained from this study. We end with a series of ideas and hints for future work.

Chapter 2

Compressed Sensing

2.1 Introduction

In this chapter we give a concise overview of the CS theory necessary for a good understanding of the following chapters. We start with the Shannon-Nyquist theorem which fixes the number of samples needed when sensing a signal. This theoretical lower bound can be decreased using Compressed Sensing, if the sensing modality and the sensed signal satisfy some conditions. These conditions on the sensed signal and on the measurement matrix, which are linked to the new sampling bound, are discussed in sections 2.2.1 and 2.2.2 of this chapter. Efficient algorithms exist for reconstructing the unknown signal from a subsampled measurement. In section 2.2.3, we discuss two families of CS reconstruction algorithms: Basis Pursuit algorithms and Matching Pursuit algorithms. We explain a couple of implementations which are used in the later chapters and compare them in terms of computation time, uniformity and robustness against subsampling.

2.2 The basic problem

2.2.1 A set of linear measurements

Suppose a discrete signal of interest x with N elements: $x = (x_1, x_2, \dots, x_N) \in \mathbb{C}^N$. In order to access this unknown vector, the signal x is sensed through a set of linear measurements. The vector of measurements $y \in \mathbb{C}^n$ is obtained after applying the measurement matrix $A \in \mathbb{C}^{n \times N}$, modelling the measurement process, on x :

$$y = Ax. \quad (2.1)$$

In order to represent a real context where measurements are noisy, a noise vector $e \in \mathbb{C}^n$ is added to the equation (2.1):

$$y = Ax + e. \quad (2.2)$$

In general, this set of n linear equations has a unique solution for x only if $n \geq N$, which bounds the minimum number of measurements n to the number of unknowns N .

2.2.2 The Shannon-Nyquist theorem

For a continuous-time signal $x(t)$, the sampling bound is determined by the Shannon-Nyquist sampling theorem. The Fourier-transform of a continuous-time signal $x(t)$ is defined by:

$$z(\omega) = \int_{-\infty}^{+\infty} x(t)e^{-j\omega t} dt. \quad (2.3)$$

x is a signal with bandwidth B if z is supported in $[-B, B]$. The Shannon-Nyquist theorem states [Jerri1977]: if $x(t)$ contains no frequencies higher than B Hertz, then it is completely determined by giving its ordinates a series of points, spaced $\frac{1}{2B}$ seconds apart. This sampling theorem dictates thus how many samples are needed to sense the continuous-time signal $x(t)$. As will be explained in detail in section 3.2.1, the discretised version of the equation 2.3 can be rewritten as the expression 2.1, where A is a Discrete Fourier Transform (DFT) matrix and where then number of rows is dictated by the Shannon-Nyquist theorem.

2.2.3 Undersampling

Classic linear algebra bounds thus the number of samples to the length of the signal. This result has a serious impact on measurement setups: (1) it sets a lower bound for the sampling rate (for example, for analogue to digital convertors), (2) if measurement times are not negligible, the number of samples will impact the total measurement time (for example, for Magnetic Resonance Imaging (MRI) and Computer Tomography (CT), where measurement times and radiation exposure must be minimized) and (3) high sampling rates can rapidly generate large data volumes which need to be transmitted and or processed (for example in Synthetic Aperture Radar (SAR) applications).

In many daily life applications, the important information is stored in only a small number of entries of the signal x compared to the total length of the signal. These signals are called sparse signals (with sparsity k) if the signal contains k non-zero elements or is called compressible if only k elements are not close to zero. Alternatively, for many non-sparse signals, the signal can be expressed in a basis in which only a small number of components are non-negligible (for example the wavelet basis for images). Since the locations of the nonzero elements in x are unknown prior to the measurement, the number of measurements (or the number of rows of A) can not be reduced or will lead to an underdetermined set of equations. This observation intuitively motivates the following questions and opens the path to sampling below the Shannon-Nyquist bound:

1. Can we sample and compress at the same time?
2. How to take compressive samples (or: What are the properties that the measurement matrix A must fulfill)?
3. How can the signal x be recovered from the compressive samples?

These questions have been examined intensively since the advent of Compressed Sensing, around 2004 and the results have been applied in a large variety of applications. The main key results, which are of importance for the later sections, are presented in the next section 2.3.

2.3 Main results in Compressed Sensing

Compressed Sensing has proven that under certain conditions on x (see: 2.3.1) and A (see: 2.3.2), a unique solution can be found from an incomplete set of linear measurements ($n < N$) in a stable and robust way and that efficient algorithms for determining this unique solution do exist.

2.3.1 Conditions on the signal

The cornerstone of the Compressed Sensing theory is the assumption that the unknown vector x is sparse, has a sparse representation or is compressible. A vector is called k -sparse if the entries of x are all zero except in a small number ($\leq k$) of coordinates:

$$x \in \mathbb{C}^N, \|x\|_0 := |\text{Supp}(x)| \leq k \leq N, \quad (2.4)$$

where $\text{Supp}(\cdot)$ denotes the support of a vector and is defined as the set of indices of the non-zero elements of x . $\|\cdot\|_0$ is a pseudo-norm indicating the sparsity of x . In general, the p -norm of a vector is defined as follows:

$$\|x\|_p := \left(\sum_{i=1}^N |x_i|^p \right)^{1/p} \quad \text{with } 1 \leq p \leq \infty \quad (2.5)$$

In practice, many signals are not exactly sparse but contain a large number of elements close to zero. These signals can be well approximated by a compressed k -sparse version of the original signal. Formally, compressible signals are defined as signals x obeying:

$$|x_i^*| \leq Ci^{-1/q}, \quad (2.6)$$

with x^* a non-increasing rearrangement of x , C is a positive constant and $0 < q < 1$. In the case that the signal is not sparse nor compressible in the measurement basis, a change of bases (Ψ) can be applied in order to obtain a sparse or compressible signal α :

$$y = Ax = A\Psi\alpha = \Theta\alpha. \quad (2.7)$$

Popular choices for Ψ are the Discrete Fourier basis, the Discrete Cosine basis and Wavelet bases which are known to be a good choice for images, etc [Lewis and Knowles1992]. The reconstruction of the vector x from an underdetermined set of equations is not self-evident since the coordinates of the nonzero elements of x are a priori unknown. The solution can be found by solving:

$$\min_x \|x\|_0 \quad \text{s.t. } y = Ax. \quad (2.8)$$

Unfortunately, this problem turns out to be NP-hard [Muthukrishnan et al.2005]. The work in [Candes et al.2006] demonstrates that the ℓ_0 -minimization can be relaxed towards an ℓ_1 -norm if the measurement matrix A satisfies certain properties (see: 2.3.2):

$$\min_x \|x\|_1 \quad \text{s.t. } y = Ax. \quad (2.9)$$

This problem is known as Basis Pursuit (BP) and is a convex minimization problem which is no longer NP-hard. $y = Ax$ only holds in an ideal noiseless world. The BP problem can be adapted for noisy signals towards:

$$\min_x \|x\|_1 \quad \text{s.t.} \quad \|Ax - y\|_2 \leq \epsilon, \quad (2.10)$$

known as Basis Pursuit Denoising (BPDN) and where $\epsilon = \|e\|_2$.

We will now summarize some important theoretical results on the properties that a the measurement matrix A must possess in order to be able to recover a k -sparse vector x , by solving (2.9) or (2.10).

2.3.2 Conditions on the measurement matrix

Null-space property

The null-space property is defined as follows [Foucart and Rauhut2013]:

A matrix $A \in \mathbb{C}^{m \times N}$ is said to satisfy the null space property relative to a set $K \subset [N]$ if:

$$\|v_k\|_1 < \|v_{\bar{k}}\|_1, \forall v \in \ker A \setminus \{0\}$$

It is said to satisfy the null space property of order k if it satisfies the null space property relative to any set K with cardinality $|K| \leq k$. $[N]$ is the ensemble $\{1, \dots, N\}$. v_k denotes a vector in \mathbb{C}^k , which is the restricted version of v to the indices in K . And $v_{\bar{k}}$ is a restricted version of v to the indices which are not in K .

Theorem [Foucart and Rauhut2013]: *Given a matrix $A \in \mathbb{C}^{n \times N}$, every k -sparse vector $x \in \mathbb{C}^N$ is the unique solution of problem (2.9) if and only if A satisfies the null space property of order k .*

Unfortunately, the null-space property is computationally hard to verify.

Coherence property

The mutual coherence of a matrix A is defined as (with a_i and a_j : the columns i and j of A):

$$\mu(A) = \max_{i \neq j} \frac{|\langle a_i, a_j \rangle|}{\|a_i\|_2 \|a_j\|_2}. \quad (2.11)$$

A sufficient condition for the recovery of all k -sparse vectors x through (2.9) is [Foucart and Rauhut2013]:

$$(2k - 1)\mu < 1. \quad (2.12)$$

The mutual coherence of a matrix A is lower bounded by [Foucart and Rauhut2013]:

$$\mu \geq \sqrt{\frac{N - n}{n(N - 1)}}. \quad (2.13)$$

Combined with the condition in (2.12), this implies that the recovery of a k -sparse vector is achievable with a number of measurements n of order k^2 . This lower bound on n can severely be improved.

Restricted Isometry Property

A popular condition on the measurement matrix is the RIP.

The Restricted Isometry Property (RIP) [Candes and Tao2005] holds with parameters (k, δ) , if:

$$(1 - \delta) \|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta) \|x\|_2^2 \quad (2.14)$$

holds for all k -sparse vectors x .

This implies that all eigenvalues of $A_K^* A_K$, where A_K is a column-submatrix of A with $K \subset [N]$ and $\text{cardinality}(K) \leq k$, are in the interval $[1 - \delta, 1 + \delta]$, or that every k columns of A are near orthogonal [Foucart and Rauhut2013]. Candès proves in [Candes et al.2006] that the problem in (2.9) is equivalent to the problem in (2.8) if A satisfies the RIP with parameters $(3k, 0.2)$. These parameters are further sharpened by Candès in [Candes2008]:

Theorem [Candes2008]: *Assume that the measurement matrix A satisfies the RIP with parameters $(2k, \sqrt{2} - 1)$. Then every k -sparse vector x can be exactly recovered from its measurements $y = Ax$ as a unique solution to the linear optimization problem (2.9).*

For measurements corrupted with noise e , Candès further proves that BPDN is stable. This means that the distance between the reconstructed vector (\hat{x}) and x is bounded:

Theorem [Candes2008]: *Let A be a measurement matrix satisfying the RIP with parameters $(2k, \sqrt{2} - 1)$. Then for any k -sparse signal x and corrupted measurements $y = Ax + e$ with $\|e\|_2 \leq \epsilon$, the solution \hat{x} to problem (2.10) satisfies: $\|\hat{x} - x\|_2 \leq C_k \epsilon$, where C_k depends only on the RIP constant δ .*

This result can be extrapolated to compressible signals:

Theorem [Candes2008]: *Let A be a measurement matrix satisfying the RIP with parameters $(2k, \sqrt{2} - 1)$. Then for any arbitrary signal x and corrupted measurements $y = Ax + e$ with $\|e\|_2 \leq \epsilon$, the solution \hat{x} to problem (2.10) satisfies: $\|\hat{x} - x\|_2 \leq C_k \epsilon + C'_k \frac{\|x - x_k\|_1}{\sqrt{k}}$, where x_k denotes the vector of the largest coefficients in magnitude of x .*

Verifying if a matrix satisfies the RIP is NP-hard [Bandeira et al.2013]. It has been proven that some matrices (for example: Gaussian, Bernoulli and partial Fourier matrices) satisfy the RIP conditions with overwhelming probability [Candes and Tao2005, Baraniuk et al.2008] [Rudelson and Vershynin2005] ("with overwhelming probability" is defined in [Candes and Tao2005] as: "with a probability decaying exponentially in n ").

2.3.3 CS reconstruction algorithms

Basis Pursuit algorithms

The BP (2.9) and BPDN (2.10) problems are convex optimization problems which can be solved by efficient algorithms. Numerous performant and competitive implementations do exist. We now briefly introduce those implementations which are used to perform the experiments in the further sections of this dissertation.

ℓ_1 -MAGIC

This is a set of recovery procedures, written by Emmanuel Candès and Justin Romberg in Matlab, for solving a convex optimization problem within 7 separate contexts. The problems are either recast as a linear program, which is the case for the BP problem and is solved using a generic path-following primal-dual method. Or, the problem is recast as a second-order cone program, solved by a log-barrier algorithm. A detailed description of both algorithms can be read in [Boyd and Vandenberghe2004]. The functions of ℓ_1 -MAGIC used for solving the BP problem: *l1_eqpd.m* and for BPDN: *l1qc_logbarrier.m* or *l1qc_newton.m*. The code, a detailed description of these functions and a users guide can be found in [Candes and Romberg2005].

SPGL1

SPGL1 is a Matlab software package, written by Ewout Van Den Berg and Michael Friedlander, for solving the BP, the BPDN and the Least Absolute Shrinkage and Selection Operator (LASSO) problem. It implements a root-finding algorithm and a spectral project-gradient algorithm. The pseudo-code and a detailed description of the algorithms are presented in [Van Den Berg and Friedlander2008] and the user manual and the SPGL1 software can be found at [SPG]. The method supposes that the noise level is approximately known. In contrast to ℓ_1 -MAGIC, SPGL1 supports complex valued coefficients.

CVX

Michael Grant and Stephen Boyd developed a solution for solving convex optimization problems, using the object-oriented features of Matlab [Grant et al.2008]. The implementation is based on two developments: (1) Disciplined Convex Programming (DCP), which is a methodology for constructing convex models. DCP consists of an atom library of functions or sets with known properties on curvature and monotonicity and of a ruleset containing a finite enumeration of ways in which atoms may be combined while preserving convexity. (2) A graph implementation which encapsulates a method for transforming instances of a specific function into a structure compatible with one of the underlying solvers. For the standard CVX user, the transformations and underlying solvers are opaque. The user writes the problem in accordance with the DCP ruleset. CVX then verifies the compliance and transforms the model to a solvable form. It then calls an appropriate solver and translates the numerical results back to the original problem. All this happens autonomously without user intervention.

Matching Pursuit algorithms

The second group of algorithms we discuss are greedy algorithms. Many variations and extensions do exist but they are mostly based on the orthogonal matching pursuit algorithm.

Orthogonal Matching Pursuit (OMP):

OMP is put forward in [Mallat and Zhang1993] as a sparse approximation algorithm for reconstructing sparse signals from an underdetermined set of measurements and the method is thoroughly analyzed by Joel Tropp and Anna Gilbert in [Tropp and Gilbert2007]. The starting point is the observation that the measurement vector y is a linear combination of k columns from A , in other words: y has a k -term representation over the matrix A . In order to reconstruct x , we need to identify the participating columns of A when taking the measurement y .

OMP is a greedy algorithm (a pseudo code for the algorithm can be found hereunder): the first iteration starts with determining which of the columns of A is the most correlated with the measurement vector y . The selected column is then added to the set of selected columns Λ_i and erased from A . The set of selected columns is used to find an new estimate for x : \hat{x} by solving a least squares problem. Finally a residual is determined by subtracting the contribution of the already selected elements of x from the measurement vector y . These steps are repeated for k iterations, hoping that the algorithm has selected and reconstructed the right k nonzero elements of x .

Algorithm 1 Orthogonal Matching Pursuit

```

1: INPUT: measurement matrix  $A(n \times N)$ ; measurement vector  $y(n \times 1)$ ; sparsity level  $k$  of
    $x(N \times 1)$ 
2: OUTPUT: estimate for  $x$ :  $\hat{x}(N \times 1)$ ; residual  $r(n \times 1)$ 
3: procedure
4:   Initialize:  $r_0 = y$ ; index set:  $\Lambda_0 = \emptyset$ ; iteration counter  $i = 1$ ;  $A_0 = \text{zeros}(n, N)$ 
5:   While  $i \leq k$  Do:
6:      $A_i = A_{i-1}$ 
7:     find index  $\lambda_i$ :  $\lambda_i = \arg \max_{j=1, \dots, N} |\langle r_{i-1}, A(:, j) \rangle|$ 
8:     augment the index set with  $\lambda_i$ :  $\Lambda_i = \Lambda_{i-1} \cup \{\lambda_i\}$ 
9:     replace the column  $\lambda_i$  of  $A_i$  by the corresponding column of  $A$ :  $A_i(:, \lambda_i) = A(:, \lambda_i)$ 
10:    set  $A(:, \lambda_i)$  equal to 0
11:    solve a least squares problem to obtain a new estimate for  $x$ :  $\hat{x} = \min_x \|y - A_i x\|_2$ 
12:    Update the residual and iteration counter:  $r_i = y - A_i \hat{x}$  and  $i = i + 1$ 
13:  End

```

The total reconstruction time for the OMP algorithm depends on the number of iterations and thus on the sparsity of the signal. For each iteration, one column of A is selected and one element of x is reconstructed. Donoho et Al. proposed a modified version of OMP, called Stagewise Orthogonal Matching Pursuit (StOMP) [Donoho et al.2012]. Instead of only selecting the largest component of the proxy $|\langle r_{i-1}, A(:, j) \rangle|$, it selects all the coordinates whose values are exceeding a chosen threshold and the algorithm iterates over an arbitrary chosen number of iterations, a value that can be chosen much lower than the sparsity of x . In their work, the authors of [Donoho et al.2012] propose to use the

notion of false alarms. A false alarm is a wrongly selected coordinate. They propose two thresholding strategies: (1) False alarm control: maximising the number of detections while guaranteeing that the total number of false alarms does not exceed $n - k$. (2) False discovery control: maximising the number of detections while guaranteeing that the number of false detections does not exceed a fixed fraction of all detections. StOMP is faster than OMP but does not provide the same reconstruction guaranties as the BP reconstruction approach.

Compressive Sampling Matching Pursuit (CoSaMP):

CoSaMP [Needell and Tropp2009] is a further improved version compared to StOMP, implementing a much larger selection of largest components of the proxy. This set of selected components is merged with the components appearing in the current approximation. The re-injection of the already selected components enables the algorithm to correct for false detections. The sparsity of the solution is assured by a pruning stage when the new approximation for x is calculated. The different steps of the algorithm can be summarized as follows [Needell and Tropp2009]:

1. Identification: A proxy of the residual is calculated from the current estimation and locates the $2k$ largest components.
2. Support merge: The set of newly detected components is merged with the components appearing in the current approximation.
3. Estimation: A least-squares problem is solved in order to obtain a new approximation for x .
4. Pruning: Only the k largest components of the elements of the solution of the least-squares problem are retained.
5. Update: A new residual is calculated.

Algorithm 2 Compressive Sampling Matching Pursuit

```

1: INPUT: measurement matrix  $A(n \times N)$ ; measurement vector  $y(n \times 1)$ ; sparsity level  $k$  of
    $x(N \times 1)$ ; maximum number of iterations  $MaxIt$ 
2: OUTPUT: estimate for  $x$ :  $\hat{x}(N \times 1)$ ; residual  $r(n \times 1)$ 
3: procedure
4:   Initialize:  $r_0 = y$ ; index set:  $\Lambda_0 = \emptyset$ ; iteration counter  $i = 1$ ;  $\text{Supp}(\hat{x}) = \emptyset$ 
5:   While  $i \leq MaxIt$  Do:
6:      $A_i = \text{zeros}(n, N)$ 
7:     find  $2k$  indices:  $\lambda_{2k} = \{\arg \max_{j=1, \dots, N} |\langle r_{i-1}, A(:, j) \rangle|\}$ 
8:      $\Lambda_i = \text{Supp}(\hat{x}) \cup \lambda_{2k}$ 
9:      $A_i(:, \Lambda_i) = A(:, \Lambda_i)$ 
10:    solve a least square problem to obtain a new estimate for  $x$ :
11:     $\hat{x}_{temp} = \min_x \|y - A_i x\|_2$ 
12:    prune the estimate, only the  $k$  largest components are retained:
13:     $\hat{x} = \max^k(\hat{x}_{temp})$ 
14:    Update the residual and iteration counter:
15:     $r_i = y - A\hat{x}$  and  $i = i + 1$ 
16:  End

```

Many other variations of the OMP algorithm have been proposed, for example: Complementary Matching Pursuit (CMP).

Matching Pursuit versus Basis Pursuit algorithms

In this subsection we make a brief comparison between the two discussed families of reconstruction algorithms: the basis pursuit algorithms, solving the BP(DN) problem and the greedy matching pursuit algorithms. The conclusions are illustrated with some experiments that are performed on simulated sparse signals sensed with a subsampled random Gaussian measurement matrix, since they are known to fulfil the conditions guaranteeing a correct BP- and MP-reconstruction. We wish a reconstruction algorithm to have following properties: a low computational complexity, a short computation time, a guarantee for a successful reconstruction, and to be robust against compressible signals and measurement noise.

1. Computation time:

Together with their simplicity to implement, greedy algorithms are mainly popular because of their short computation times compared to BP algorithms under certain settings. This is illustrated in Fig 2.1 and Fig 2.2. For this experiment 64 random sparse signals are created (the positions of the non-zero elements are randomly picked and the values are drawn from a standard normal distribution) with signal length (N) equal to 512 and a sparsity rate (k/N) of 5%. These signals are measured through a random Gaussian matrix. The measurements are subsampled with a subsampling rate (n/N) of 30%. Fig 2.1 depicts the mean reconstruction times over the 64 executions. The experiment is repeated for increasing dimensions of N from 256 up to 1600. We observe that for small signal lengths (< 600) both OMP and CoSaMP outperform all BP algorithms. We further remark that the CoSaMP is largely faster than all the other algorithms. Fig 2.2 then, shows the results for a fixed signal length equal to 512 and a varying sparsity rate between 2% and 50%. We can again conclude that CoSaMP has the shortest computation time under all circumstances. We also remark the relative fast reconstruction times of SPGL1 compared to ℓ_1 -MAGIC and CVX in both experiments.

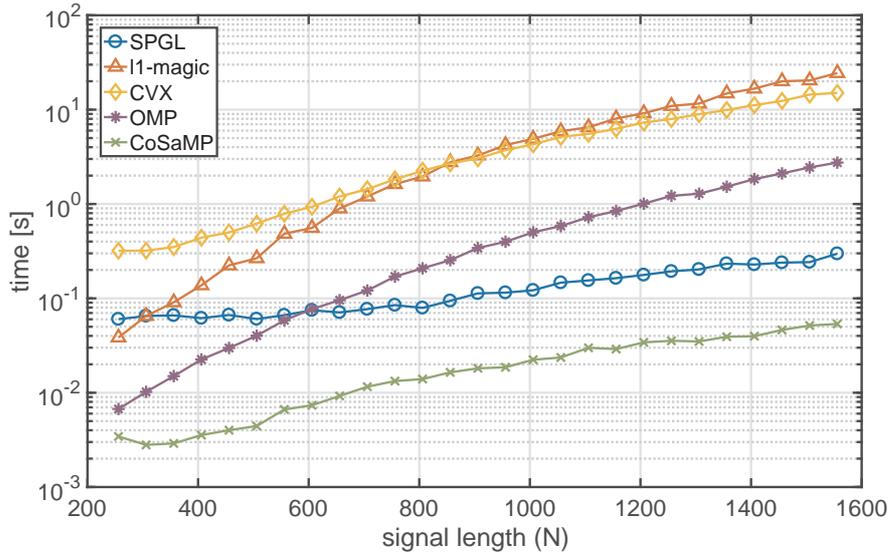


FIGURE 2.1: Mean reconstruction time (over 64 iterations) for increasing signal lengths (N) for random vectors ($k/N = 5\%$ and $n/N = 30\%$). The reconstructions are obtained with (1) the BP algorithms: SPGL1, ℓ_1 -magic and CVX and (2) the MP algorithms: OMP and COSAMP.

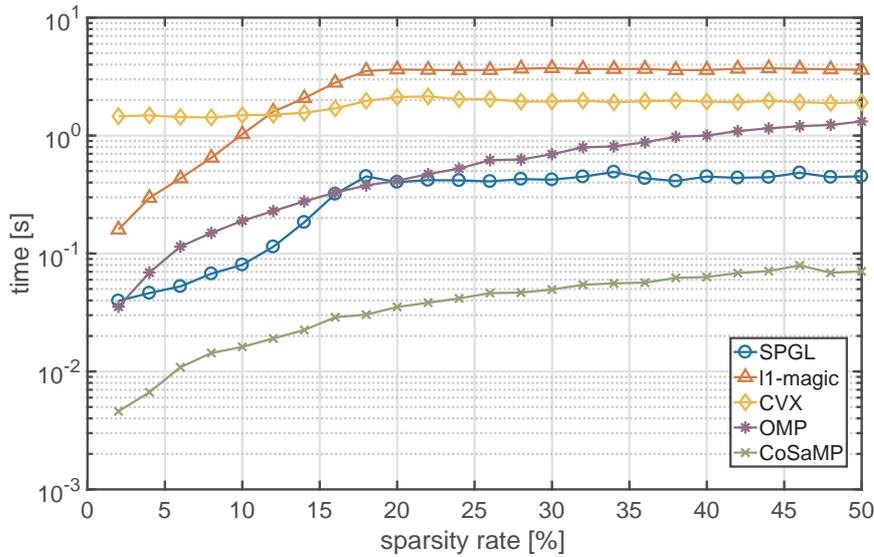


FIGURE 2.2: Mean reconstruction time (over 64 iterations) for increasing sparsity rates (k/N) for random vectors ($N = 512$ and $n/N = 30\%$). The reconstructions are obtained with (1) the BP algorithms: SPGL1, ℓ_1 -magic and CVX and (2) the MP algorithms: OMP and COSAMP.

2. Uniformity:

The recovery by an algorithm is defined to be uniform [Lin and Li2013] if, with high probability, the support of every sparse signal can be reconstructed under appropriate conditions. Whereas nonuniform recovery guarantees that for a given sparse signal, its support can be reconstructed on the draw of the random matrix, with high probability and under appropriate conditions. It is reported, that in contrast with the BP algorithms, OMP is nonuniform [Needell and Tropp2009]. CoSaMP on the other hand guarantees uniformity [Needell and Tropp2009]. We illustrate this by performing the following Monte Carlo test:

- First we create 3000 random vectors of length N equal to 256 and with only $k = 8$ nonzero elements which are randomly selected.
- Each of these vectors are sensed by a random Gaussian matrix of size 77×256 .
- Using these measurements, the original signals are reconstructed with three different CS reconstruction algorithms: a basis pursuit algorithm (CVX), OMP and CoSaMP.
- In order to assess the reconstruction performance, the relative error is measured after each reconstruction. The relative error is the ratio of the Euclidean distance between the reconstructed and the original signal to the ℓ_2 -norm of the original signal. The results are depicted, using boxplots, in Figure 2.3 (a) for the reconstruction of the noiseless vectors.
- Figure 2.3 (b) shows the results obtained after performing the same experiment after adding white gaussian noise, with a signal to noise ratio equal to 15 dB, to the measurements.

We observe in Figure 2.3, that with all of the three reconstruction algorithms low relative errors are obtained. We further see that only the OMP fails to deliver a good reconstruction for all the vectors. This nonuniformity property of OMP was summarised as follows by Tropp et al. in [Tropp and Gilbert2007]: With Gaussian or Bernoulli measurement matrix, BP can, with high probability, recover all sparse signals. In the same setting, OMP recovers each sparse signal with high probability but with high probability fails to recover all sparse signals.

In figure 2.3 (b), we discover the excellent reconstruction from the subsampled measurements contaminated with noise, when using a BPDN algorithm (CVX) and if the noise level is well estimated.

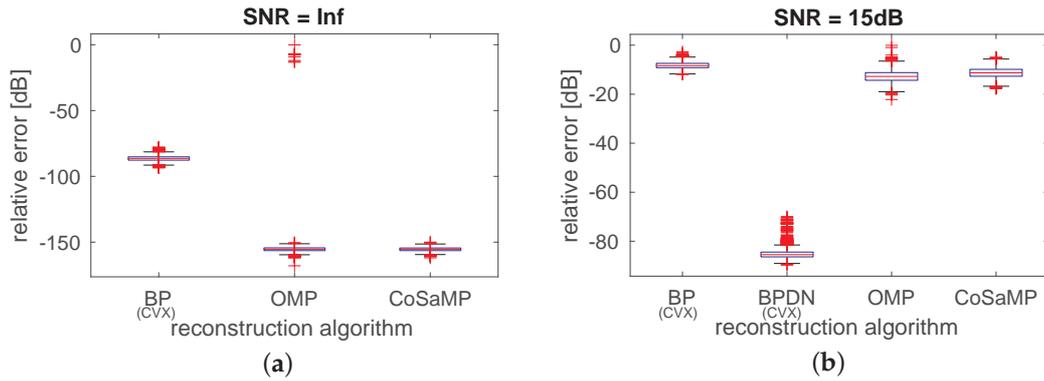


FIGURE 2.3: Boxplots of the relative error for BP (CVX) and MP (OMP and CoSaMP) algorithms. In (a) : reconstruction from noiseless measurements and in (b): reconstruction from measurements with a SNR = 15 dB.

3. Robustness:

The robustness of the CS algorithms against subsampling is, besides the computation time and the uniformity, a very important property for evaluating the CS algorithm. From the theorems we discussed in this chapter, it is clear that the reconstruction performance is a function of the sparsity of the signal to reconstruct. In order to evaluate the robustness of the reconstruction algorithms, we perform a series of tests to measure the mean relative reconstruction error for the possible combinations of N , n and k . The relative error is the ratio of the Euclidean distance between the reconstructed and the original signal to the ℓ_2 -norm of the original signal. These results are presented in phase diagrams, which were introduced by Donoho and Tanner in [Donoho and Tanner2009]. The mean relative reconstruction error is calculated for all combinations of the undersampling fraction $\delta = n/N$ (the ratio of the number of samples to the signal length) and the sparsity ratio $\rho = k/n$ (the ratio of the number of nonzero elements in the original signal to the number of samples). The results of these tests are represented in a diagram with on the horizontal axis the undersampling fraction and on the vertical axis the sparsity ratio.

The diagrams in Figure 2.4 depict several phase diagrams. The diagrams on the left hand side represent the relative errors calculated from 256 reconstructions of each combination of ρ and δ . We performed this experiment with three different CS reconstruction algorithms: (1) a basis pursuit algorithm: SPGL1, (2) OMP and (3) CoSaMP. We observe that the CoSaMP algorithm has the smallest blue zone, which characterises the combinations of ρ and δ delivering a very low relative reconstruction error. The diagrams on the right hand side depict the mean reconstruction time spent by the three different CS algorithms to obtain a reconstruction for the different combinations of ρ and δ . As expected, the basis pursuit algorithm is the slowest algorithm. Whereas CoSaMP is many times faster than the other algorithms for all possible combinations of ρ and δ .

The results presented and discussed in this chapter were obtained by performing synthetic measurements with a random Gaussian matrix. The matrices modelling the sensors utilized in chapters 4 and 5, are the Discrete Fourier matrix and the Synthetic Aperture

Radar measurement matrix. In the next chapter we will focus on the performance of CS in combination with these matrices.

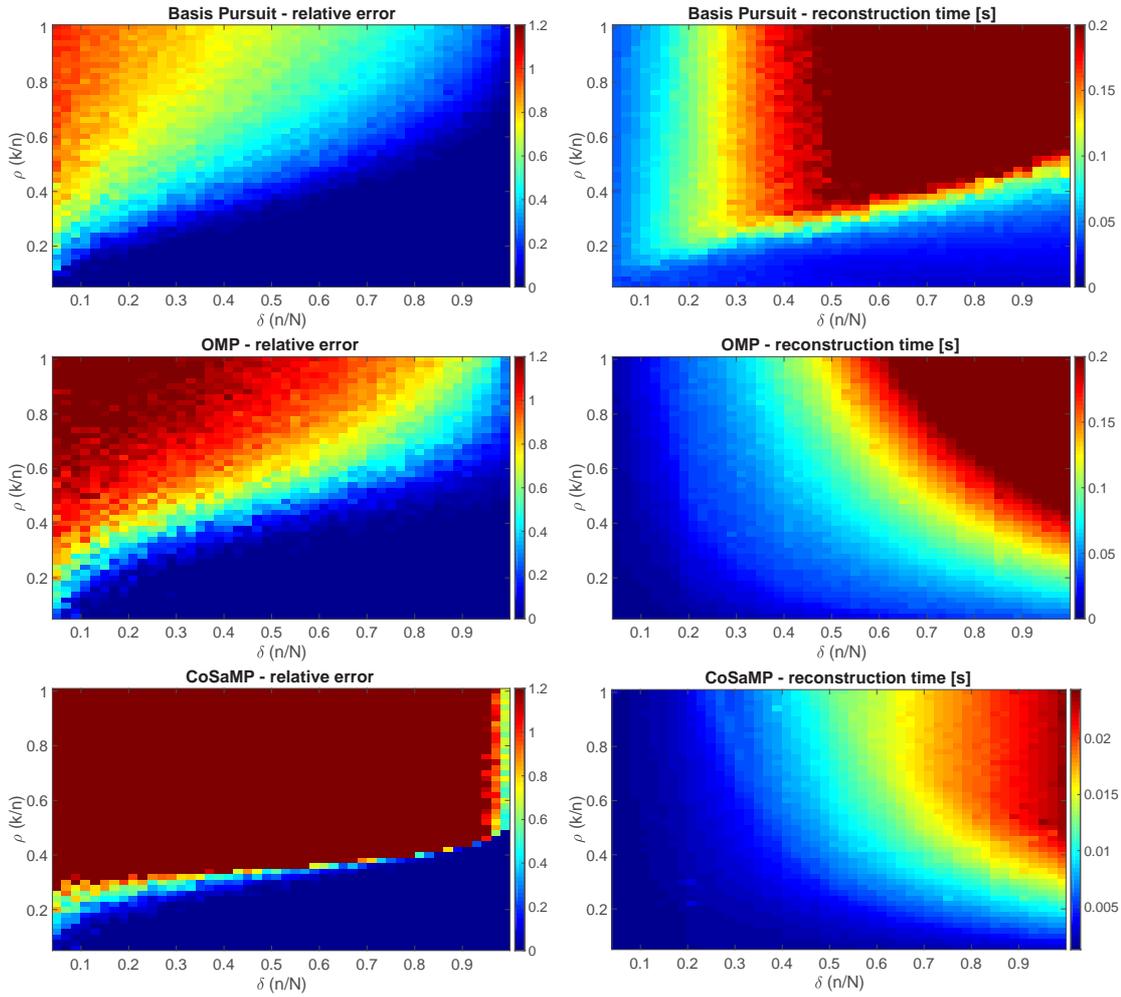


FIGURE 2.4: Phase diagrams using a BP algorithm (SPGL1), OMP and CoSaMP. On the left hand side: the mean relative reconstruction error as a function of the undersampling fraction δ and the sparsity ratio ρ . On the right hand side: the computation time as a function of δ and ρ .

Chapter 3

Compressed Sensing applied on SFCW and SAR measurements

3.1 Introduction

In this chapter we concentrate on using CS together with two different measurement modalities: (1) Stepped-Frequency Continuous Wave (SFCW) measurements and (2) Synthetic Aperture Radar (SAR) measurements. These two modalities will be used for obtaining the real data for the experiments described in the chapters 4 and 5. We will start with the description of the Discrete Fourier Transform (DFT) and we will link this transform to the SFCW measurement matrix. Through a series of tests, presented in sections 2 and 3, we will show that Compressed Sensing can lower the demanding upper bound for the number of samples dictated by the Shannon-Nyquist theorem for as well SFCW as for the SAR measurements. This lower bound is a function of the sparsity of the signal, which is a priori unknown for most applications. In other words, we know with high probability that we can obtain an exact reconstruction of the sensed signal from a much lower number of samples, but, unfortunately, we can not tell a priori how many samples we have to take. A solution to this problem, named sequential CS, is explored in section 4. In section 5 we describe a series of signal processing methods, such as the Singular Value Decomposition (SVD) and the Principal Component Analysis (PCA), which can be combined with CS to obtain an even better reconstruction performance with CS under certain circumstances. The corner stone of CS is the prior assumption that the signal is sparse or has a sparse representation in a known basis. This prior knowledge is sometimes not the only knowledge we have on the signal before starting the measurement. In section 6 we will test and evaluate different methods for adding this supplementary information, known as Side Information (SI), into the CS reconstruction. We will treat two distinct cases: (1) Homogeneous SI: where the SI is obtained with the same sensing modality and (2) Heterogeneous SI: where the SI is obtained through a different measurement system. This chapter can be perceived as a toolbox of methods which we explore and evaluate on synthetic data in this chapter and which will be used on real data, when appropriate, in the chapters 4 and 5 of this dissertation.

3.2 SFCW measurements

3.2.1 The Discrete Fourier Transform

The Discrete Fourier Transform (DFT) is the discretised version of the continuous Fourier transform $z(\omega)$ of a continuous function $x(t)$:

$$z(\omega) = \int_{-\infty}^{+\infty} x(t)e^{-j\omega t} dt, \quad (3.1)$$

where the function $x(t)$ is known in N equidistant samples, separated by a time interval T :

$$z(\omega) = \sum_{l=0}^{N-1} x(l)e^{-j\omega lT}. \quad (3.2)$$

Since there is only a finite number of samples, the DFT treats the data as if they were periodic (with a period equal to NT). The equation (3.2) can thus be rewritten as:

$$z(n) = \sum_{l=0}^{N-1} x(l)e^{-j\frac{2\pi}{N}nl} \quad \text{with } n \in [0 : N - 1]. \quad (3.3)$$

Or expressed in matrix form:

$$\begin{bmatrix} z(0) \\ z(1) \\ z(2) \\ \vdots \\ z(N-1) \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & W & W^2 & \dots & W^{N-1} \\ 1 & W^2 & W^4 & \dots & W^{2(N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & W^{N-1} & W^{2(N-1)} & \dots & W^{(N-1)(N-1)} \end{bmatrix} \begin{bmatrix} x(0) \\ x(1) \\ x(2) \\ \vdots \\ x(N-1) \end{bmatrix}, \quad (3.4)$$

with $W = e^{-j2\pi/N}$.

The minimum number of samples N , needed to capture all the information in $x(t)$, is dictated by the Shannon-Nyquist theorem. We will now demonstrate the aliasing problem that appears when sampling at a sampling rate below the Nyquist sampling rate and we show how to break through this bound with Compressed Sensing.

Let us consider the following signal:

$$x(t) = 3 + 6 \sin(2\pi 550t) + 4 \sin(2\pi 600t) \quad (3.5)$$

This signal, depicted by the blue curve over a time interval of 20 ms in Figure 3.1, has one DC component, one component at a frequency of 550 Hz and a component at 600 Hz. This implies that the signal should be sampled at a sampling frequency of at least 1200 Hz. Let us now sample this signal at two different sampling frequencies: (1) at 2 kHz, which is largely above the Nyquist sampling rate (the samples in the represented time interval are encircled in red in Figure 3.1) and (2) at 1 kHz, which is below the Nyquist sampling bound (highlighted with the yellow symbols). The second graph in Figure 3.1 shows the spectrum of the signal sampled at 2 kHz. As expected, we obtain three peaks:

the first peak representing the DC component at 0 Hz and the two components at 550 Hz and 600 Hz. If we now look at the spectrum of the signal sampled at 1 kHz (the third graph of Figure 3.1), two frequencies appear: the frequencies 550 Hz and 600 Hz are aliased into the frequencies 450 Hz and 400 Hz respectively, because the signal is copied in integer multiples of 1000 Hz on both sides of the frequency axis.

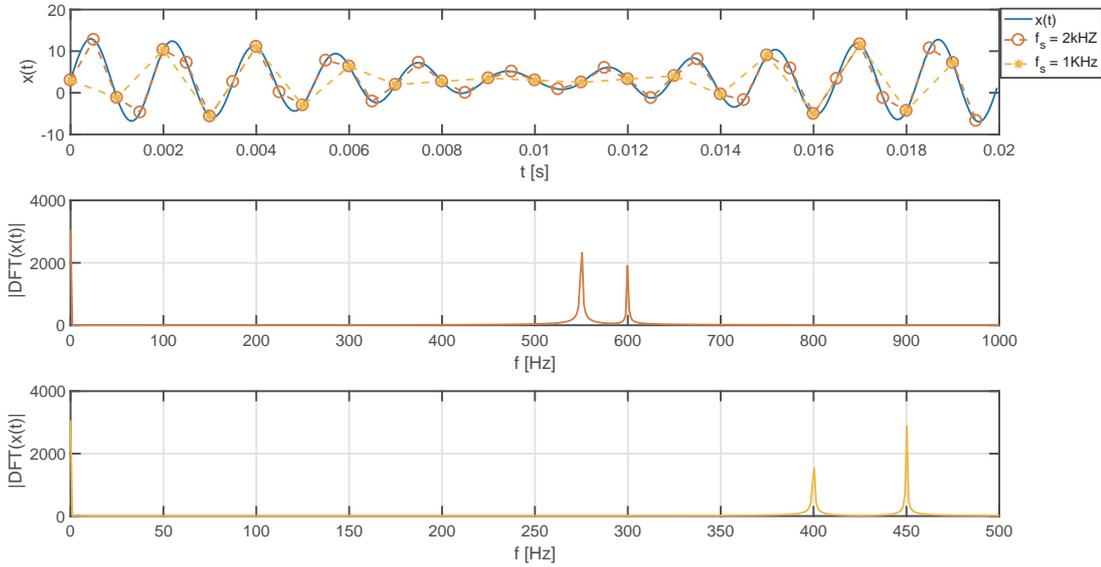


FIGURE 3.1: The upper graph: the original signal (in blue) with in red: the selected samples at a sampling rate of 2 kHz and in yellow the samples selected at a sampling rate equal to 1 kHz. The middle graph: the spectrum of the signal sampled at 2 kHz. The lower graph: the spectrum of the signal sampled at 1 kHz.

Let us now consider the same signal over a time interval of 2 seconds. The minimum number of samples, needed at the Nyquist rate, is thus equal to 2400. From the original samples, equidistantly taken over the 2 seconds time interval at the sampling rate of 2 kHz, we will now select only 2000 samples using two different strategies: (1) we take 2000 equidistant samples and (2) we randomly select 2000 samples out of the 4000 original samples. From these undersampled signals, we will try to reconstruct the original 4000 samples using a CS algorithm (SPGL1) by exploiting the fact that the signal is supposed to be sparse in the frequency domain. In other words, we will solve the following minimization problem:

$$\min \|z\|_1 \quad \text{s.t.} \quad y = S\Psi^{-1}z, \tag{3.6}$$

where Ψ is a DFT matrix and serves as a sparsifying transform, S is a matrix of size $n \times N$ which is obtained by randomly selecting n rows from the unit matrix and z is the DFT-transform of x . y corresponds thus to the subsampled signal x . The first graph in Figure 3.2 shows again the original signal (3.5) in blue, the equidistant selected samples

are encircled in red and the randomly selected samples are marked with the yellow symbols. The second and the third graph depict the spectrum of the reconstructed signals from respectively, the 2000 equidistant samples and from the randomly selected samples. We already note the importance of the subsampling strategy: in both cases, the signals are reconstructed from 2000 samples, but the signal is only correctly reconstructed, thus without aliasing, with the randomly selected samples.

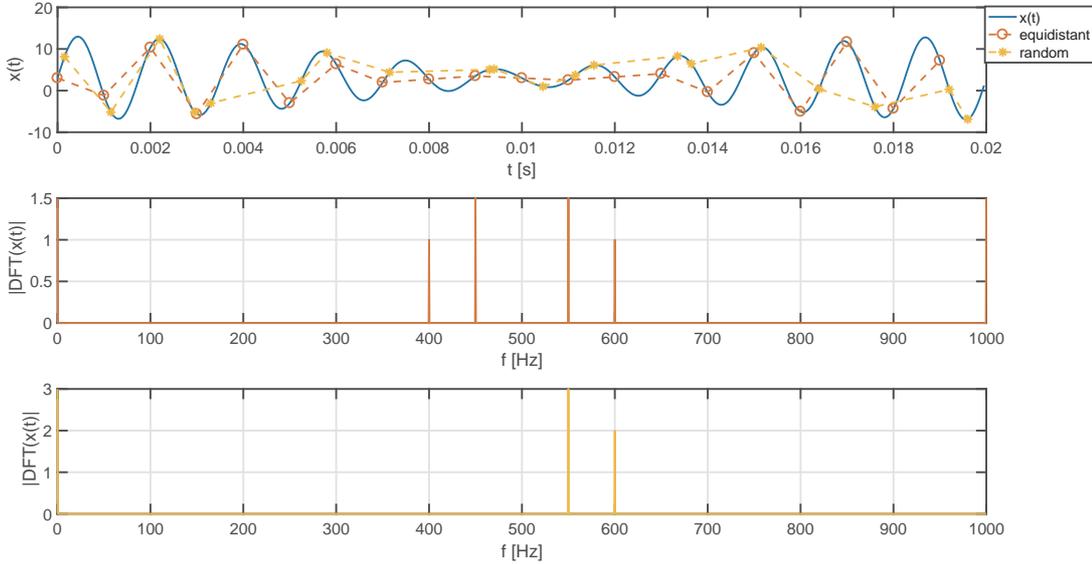


FIGURE 3.2: The upper graph: the original signal (in blue) with marked in red: the equidistant samples and in yellow: the randomly selected samples. The middle graph: the spectrum of the reconstructed signal from the equidistant samples. The lower graph: the spectrum of the reconstructed signal from the randomly selected samples.

3.2.2 Stepped-Frequency Continuous Wave measurements

We will now focus on Stepped-Frequency Continuous Wave (SFCW) measurements. A SFCW sensor, as will be used during the real experiments discussed in chapter 4, emits L discrete frequencies f_l ($l = 0, \dots, L - 1$) evenly spaced (with a frequency step equal to Δ_f) over the total bandwidth B . The signal is thus described by the following equation:

$$s(t) = \sum_{l=0}^{L-1} \text{rect}\left(\frac{t - lT_p}{T_p}\right) \exp(j2\pi f_l t), \quad (3.7)$$

where $\text{rect}(\cdot)$ is the rectangular function and T_p is the period during which the total bandwidth B is emitted (Figure 3.3).

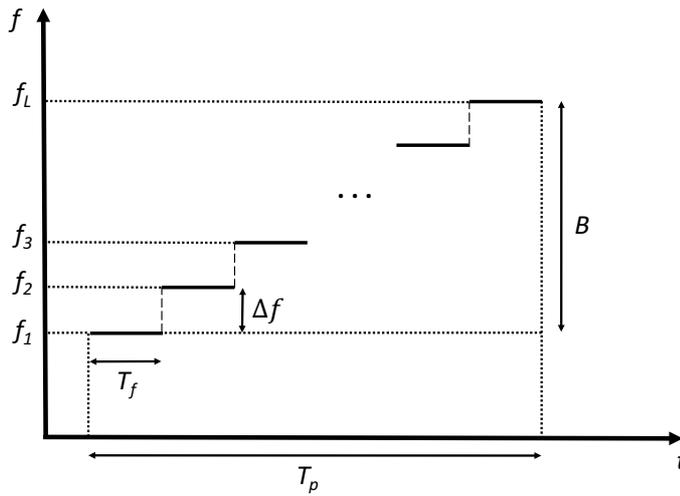


FIGURE 3.3: Stepped-frequency continuous wave signal.

The received signal is the sum of all K backscattered signals coming from K different reflectors. Each of the K signals are received by the sensor antenna after a time equal to the Round Trip Time (τ). This τ is equal to $\frac{2R_k}{c}$ where R_k is the distance between the reflector k and the antenna and c is the propagation speed of the electromagnetic waves through the medium ($c = \frac{c_0}{\sqrt{\mu\epsilon}}$). The total received signal is thus:

$$s(t) = \sum_{k=0}^{K-1} \sum_{l=0}^{L-1} a_{k,l} \exp\{j2\pi f_l(t - \tau_k)\}. \quad (3.8)$$

Which, after homodyne demodulation, becomes:

$$s_b(t) = \sum_{k=0}^{K-1} \sum_{l=0}^{L-1} a_{k,l} \exp\{-j2\pi f_l \tau_k\}. \quad (3.9)$$

This beat signal is the DFT of the reflectivity function $x = \sum_{k=0}^{K-1} \sum_{l=0}^{L-1} a_{k,l}$. As was explained earlier, if we want to avoid aliasing in the reflectivity function, the sampling rate in baseband, should at least be equal to $2B$. Suppose we want to avoid aliasing over a range equal to R_{max} , which corresponds to a τ equal to $\frac{2R_{max}}{c}$, the frequency step size should thus at most be equal to:

$$\Delta f = \frac{c}{2R_{max}}. \quad (3.10)$$

Let us illustrate this with an example.

We synthetically create an SFCW sensor with following parameters:

- Frequency bandwidth: 8 GHz - 12 GHz
- Range resolution: $\frac{c}{2B} = 0.0375m$

With this virtual sensor we want to sense the reflectivity function over a distance of 7.5 m. This maximum distance dictates that, according to the equation (3.10), Δ_f should be equal to 20 MHz. In other words, at least 200 samples are needed if we want to avoid aliasing. To test this number, we create a random reflectivity function containing 5 random nonzero elements and we sense this signal by multiplying it with the DFT-matrix. Now that we have the measurement at the Nyquist rate, we try to reconstruct the original reflectivity function from a decreasing number of samples. The reconstructions are obtained by performing an Inverse Fast Fourier Transform (IFFT) of the sampled measurement and are then evaluated by calculating the Euclidean distance between the original reflectivity vector and the corresponding reconstructed vector. The blue graph in Figure 3.4 depicts the reconstruction error obtained after the reconstruction from 4 samples up to 200 samples. As expected, 200 samples are necessary in order to obtain a correct reconstruction of the reflectivity function.

3.2.3 CS applied on SFCW measurements

We now repeat the same experiment, but this time, we use a CS algorithm (SPGL-1) to reconstruct the reflectivity function by solving the following minimization problem:

$$\min \|x\|_1 \quad \text{s.t.} \quad y = Ax, \quad (3.11)$$

where y is the randomly sampled SFCW measurement (n samples), A is the DFT matrix ($n \times N$) and N is the length of the original reflectivity function. We observe that the reconstruction error stays equal to zero when decreasing the number of samples down to only 40. This result illustrates that, thanks to CS, the number of samples can be reduced severely below the number dictated by the Shannon-Nyquist theorem, for SFCW measurements, if the reflectivity function happens to be sparse.

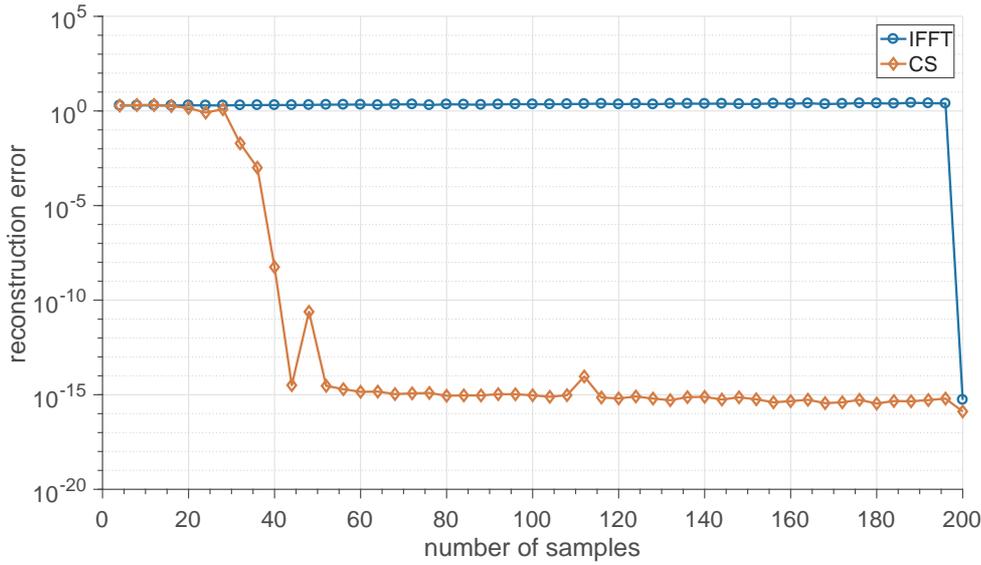


FIGURE 3.4: Reconstruction error obtained after the reconstruction of a signal from a sampled SFCW measurement, for an increasing number of samples, by applying an IFFT (in blue) or a CS reconstruction algorithm (SPGL1) (in red).

The quality of the reconstructed reflectivity function does not only depend on the number of samples and the sparsity of the scene. The reconstruction performance using a DFT matrix depends also on:

1. The CS reconstruction algorithm
2. The adopted sampling strategy

CS reconstruction evaluation

In order to obtain a clearer view on the CS performance for SFCW measurements, we will now conduct a series of tests which will result in phase diagrams, depicting the probability of a good reconstruction estimated from 100 trials and this for each possible combination between the undersampling fraction $\delta = n/N$ (number of samples/length of the reflectivity function) and the sparsity ratio $\rho = k/n$ (number of nonzero elements of the reflectivity function/number of samples). The representation of the results by a probability of successful reconstruction as a function of the undersampling rate and sparsity ratio was first suggested in [Donoho and Tanner2009] and has the advantage to make the reconstruction performances easy to interpret and compare. For this set of measurements N is chosen to be equal to 200 and a reconstruction is categorised as successful if the relative error is smaller than or equal to 10^{-3} , which is an arbitrarily chosen bound. When decreasing the sparsity ratio the relative error decreases steeply to extremely low values. This behavior limits the impact of the choice of the bound:

$$\frac{\|\hat{x} - x\|_2}{\|x\|_2} \leq 10^{-3}, \tag{3.12}$$

where x is the original reflectivity function and \hat{x} is the CS reconstructed reflectivity function. The probability of successful reconstruction is calculated over 100 trials and the same test is repeated three times. Each time, a different CS reconstruction algorithm is tested:

1. The results obtained with the SPGL-1 Basis Pursuit algorithm are depicted in Figure 3.5. We can clearly identify two regions (or phases) in the diagram: (1) The red area, where the probability to obtain a good reconstruction is equal to 100% and, (2) the blue area where the probability is equal to 0%. We further observe that the transition between the two phases is abrupt. This means that the zone, where the success of the reconstruction for a combination of N , n and k is uncertain, is small.

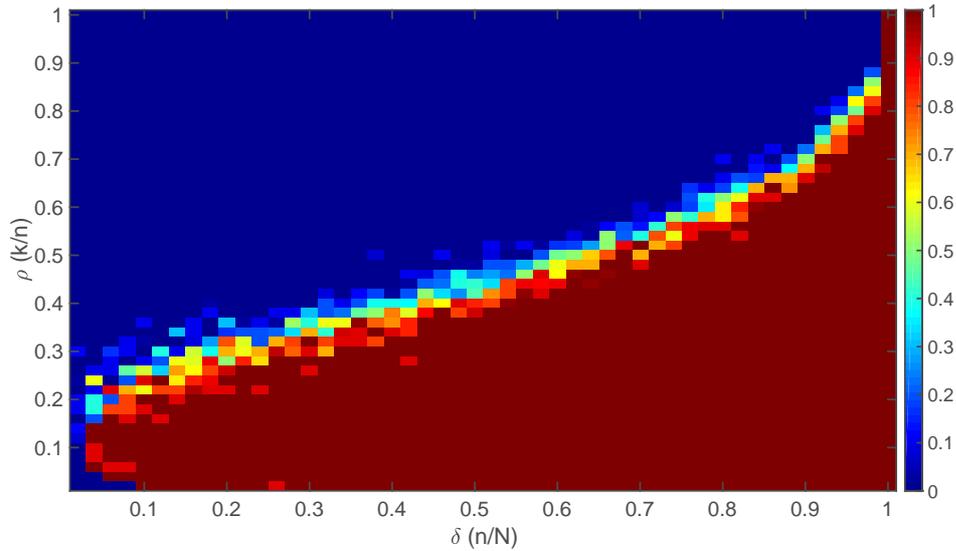


FIGURE 3.5: Phase diagram obtained using SFCW measurements with SPGL-1. The color scheme depicts the probability of a successful reconstruction. We observe two distinct areas: the blue area where the probability is equal to 0 and the red area where the probability is equal to 1.

2. The phase diagram obtained with CVX is shown in figure 3.6. We see the same trend as for the results obtained from the SPGL-1 algorithm. The red and blue areas are comparable but the transition zone is even smaller with CVX.

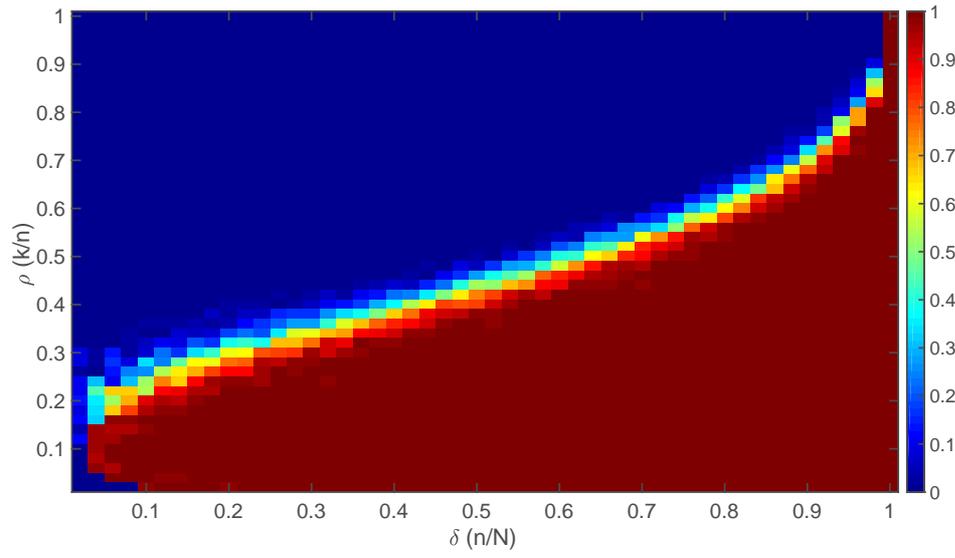


FIGURE 3.6: Phase diagram obtained using SFCW measurements with CVX. The color scheme depicts the probability of a successful reconstruction. (blue: probability = 0; red: probability = 1).

3. Figure 3.7 shows the results obtained with a matching pursuit algorithm: OMP. First, we observe that the red area is large and thus that this algorithm is successful over a wide range of combinations of N , n and k . But, on the other hand we also see a large transition zone, where the outcome of the reconstruction is unsure. Finally, to be honest, we have to add the comment that we chose the number of iterations exactly equal to the number of nonzero elements k of the signal to reconstruct. This number k is, in most applications, unknown and must be guessed and the reconstruction quality can thus deteriorate due to reconstruction noise.

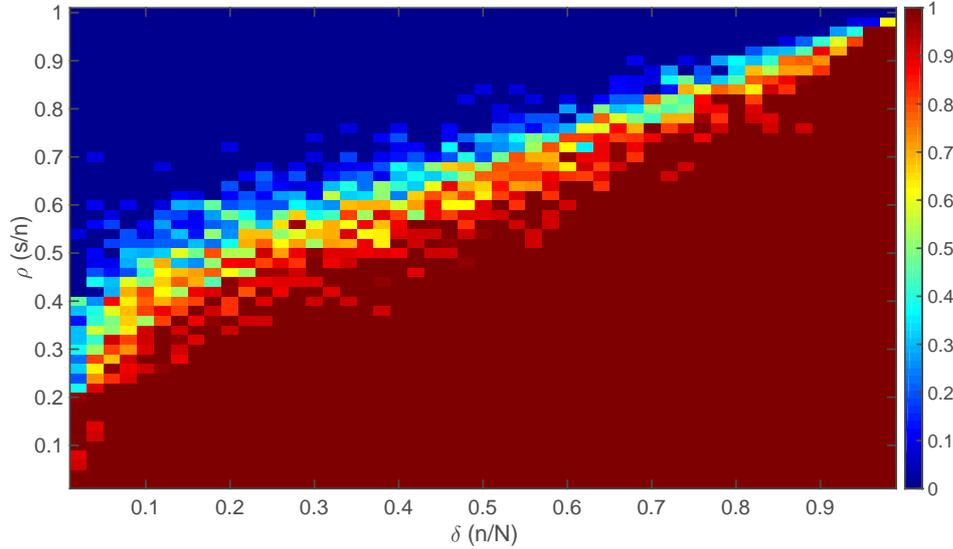


FIGURE 3.7: Phase diagram obtained using SFCW measurements with OMP. The color scheme depicts the probability of a successful reconstruction. (blue: probability = 0; red: probability = 1).

Sampling strategies

Compressed Sensing theory does not put restrictions on the way we have to sample when using subgaussian matrices, allowing to take equidistant samples as we did in Chapter 2. But, as we discovered during the tests performed on CS with the DFT (3.2), the sampling strategy can have a serious impact on the reconstruction performance when using SFCW measurements. In general, a random sampling strategy is known, for example in Magnetic Resonance Imagery (MRI) [Haldar et al.2011, Duarte and Eldar2011], to produce good results. The work in [Xu and Xu2015] proposes a deterministic selection of the samples, or the rows of the DFT matrix, using Katz' character sum estimation [Katz1989]. We now use the same experimental parameters as in [Xu and Xu2015] for comparing: (1) the random sampling scheme, (2) the deterministic approach and, (3) the equidistant sampling strategy. Therefore, we create random vectors of length $N = 840$, with an increasing number of nonzero elements ($k/N = 0.02, 0.04, \dots, 0.2$). These vectors are then synthetically measured by applying the DFT-matrix in which only 29 rows are randomly selected using the three different sampling strategies. This test is repeated 200 times and

the reconstructions (obtained with CVX) are evaluated by calculating the reconstruction success rate (a reconstruction is considered successful if the Euclidean distance between the original vector and the reconstructed vector is smaller or equal to 0.001). The results of this experiment are depicted in Figure 3.8 together with the symmetric confidence intervals ($= \text{success rate} \pm 1.96\sqrt{\frac{0.25}{200}}$). We clearly see that the equidistant sampling is not a good choice to subsample the SFCW measurements. We also observe that the random and deterministic selection perform equally. Contrary to what is mentioned in [Xu and Xu2015], the deterministic approach does not add a significant improvement to the reconstruction performance compared to the random sampling scheme.

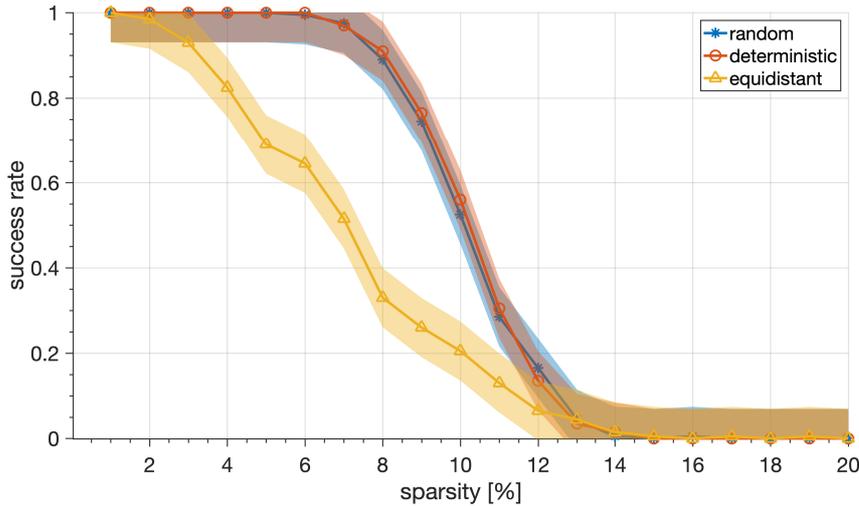


FIGURE 3.8: Reconstruction success rate (with the confidence intervals) for different sampling strategies ($N = 840$, and $n = 29$).

3.3 SAR measurements

Synthetic Aperture Radar (SAR) measurements can somehow be seen as a 2-D extension of the 1-D Frequency Modulated Continuous Wave (FMCW) measurements. In this work, we will restrict ourselves to SAR systems based on a Vector Network Analyser (VNA), measuring the S_{11} reflection parameter at a number of discrete frequencies within a chosen bandwidth B :

$$S_{11}(f_p) = \frac{S_{Rx}(f_p)}{S_E(f_p)}, \quad (3.13)$$

where $S_{Rx}(f_p)$ and $S_E(f_p)$ are the complex amplitudes of the received and the emitted signal at frequency f_p respectively. In other words this type of SAR emit SFCW signals. In our work, we suppose that the Vector Network Analyser (VNA) emits a series of independent signals at discrete frequencies. In reality, the signals have a bandwidth, defined as the Resolution Bandwidth (RBW), since an infinite small bandwidth would result in an infinite long sweep time. A smaller RBW lowers the noise floor, but extends thus the

sweep time over the chosen span of frequencies. Thanks to a large total bandwidth B , a reflectivity function with an excellent resolution can be obtained, equal to $\frac{c}{2B}$, in the dimension of the propagation direction of the electromagnetic waves which is called the range direction. The resolution in the cross-range direction however, is determined by the dimensions of the antenna: the larger the antenna, the smaller the antenna beam width, and thus the better the resolution in the cross-range dimension. In order to obtain good quality images, the resolutions in the range and cross-range dimensions should be within the same order of magnitude. This implies that large antennas are required. The SAR offers a solution to this problem and avoids the necessity for an inconveniently large antenna. The required large aperture is created synthetically by moving the antenna, which can have small physical dimensions, in the cross-range direction (Figure 3.9). The resolution obtained with a SAR in the cross-range dimension can be approximated by [Oliver and Quegan2004]:

$$\Delta r_{cross-range} = \frac{c}{2f_c \sin \theta'} \quad (3.14)$$

where f_c is the center frequency of the total bandwidth and θ is the opening (-3dB) of the real antenna. The raw data obtained from this type of data is stored in a raw data matrix where the row t contains the complete measurement obtained at the sensor position sp_t and the column p contains the S_{11} coefficient at frequency f_p over all the sensor positions.

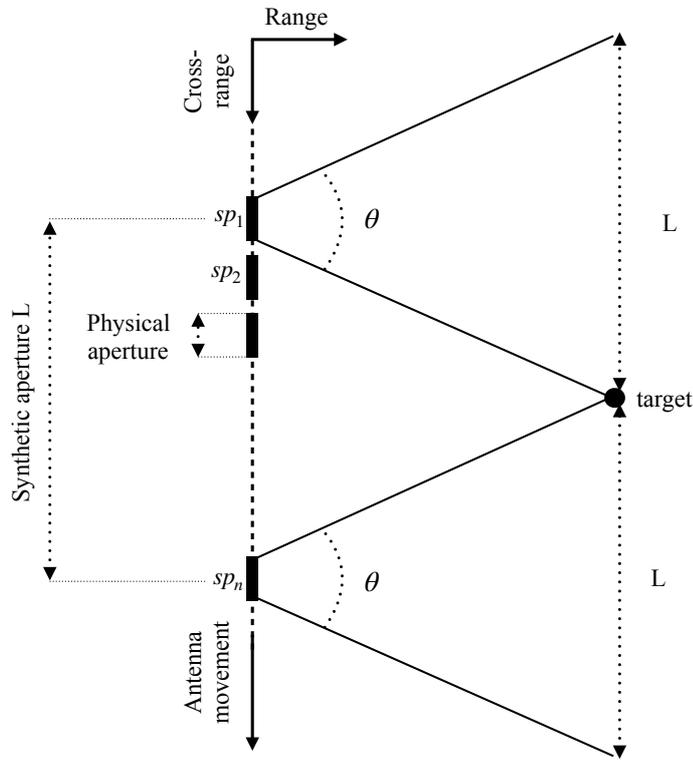


FIGURE 3.9: Schematic representation of the working principal of a SAR with physical opening angle θ and synthetic aperture L .

The raw data needs to be compressed in both the range and azimuth dimension in order to obtain the 2-D reflectivity function. Since the signals in the range dimension are SFCW measurements, this compression can be performed by a simple Inverse Fast Fourier Transform (IFFT) of each row individually. Several well-known SAR compression algorithms do exist to perform a compression in the cross-range dimension, for example: Time domain algorithms [Ribalta2011], Range-doppler SAR processing algorithms [de Wit et al.2006] or wavenumber domain focusing algorithms [Bamler1992].

Alternatively, we will use Compressed Sensing to construct the SAR image from the raw data. The reason for using CS for the SAR compression (in range and cross-range) is that CS allows us to obtain a correct SAR image from an incomplete measurement, under the condition that the reflectivity function happens to be sparse. The SAR image is obtained by solving following minimization equation:

$$\min \|x\|_1 \quad \text{s.t.} \quad y = Ax, \quad (3.15)$$

where y is the incomplete SAR measurement, A is the SAR measurement matrix which maps the vector x , containing the reflectivity function of the sensed scene, to the SAR measurement.

The SAR measurement matrix

We will now explain how we build the measurement matrix that models the SAR measurement [Wei et al.2010]. It is clear that a well-modeled measurement matrix A is needed to obtain a good reconstruction. The matrix A is a complex matrix of size $n \times N$, where N is the number of pixels composing the SAR image and n is the size of the incomplete SAR measurement. Each column can thus be attributed to a single pixel of the SAR image and each row is linked to a single sample of the SAR measurement.

The SAR measurement matrix is column-wise populated with the vectorized sensor response to a reflectivity function containing a single reflective component. This means that we assume that the scene is populated with point targets, which is not the case in reality. We concentrate now on the column corresponding to the pixel x_{ij} . This column contains thus a synthetic SAR measurement of a scene of dimensions n_f (equal to the number of frequencies) in range and n_{sp} (which corresponds to the number of sensor positions) in cross-range. This scene is characterised by a reflectivity function equal to zero except in only one element, corresponding to the position x_{ij} (Figure 3.10). This column will thus contain the synthetically created S_{11} coefficients obtained for all the combinations of the f_n frequency and the sp_n sensor positions (3.16). The SAR measurement matrix is thus populated with measurements of scenes containing a single point target. This has a consequence that the CS reconstruction is performed supposing that the reflectivity function is the sum of point targets (which is not the case in reality).

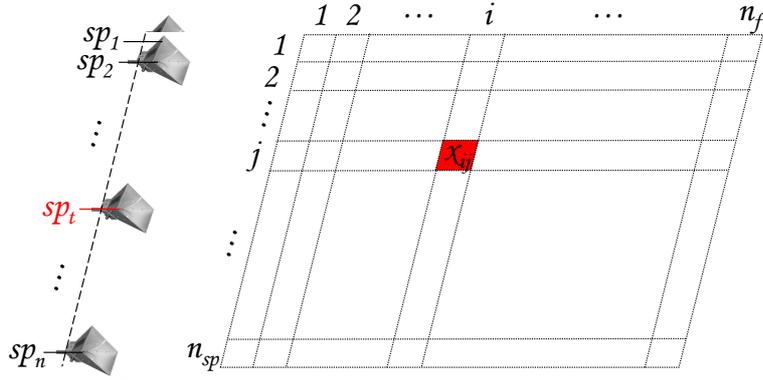


FIGURE 3.10: Schematic representation of the scene sensed to fill the columns of A corresponding to the pixel x_{ij} .

$$A = \begin{bmatrix} S_{11}^{x_{1,1}}(sp_1, f_1) & S_{11}^{x_{1,2}}(sp_1, f_1) & \dots & S_{11}^{x_{2,1}}(sp_1, f_1) & \dots & S_{11}^{x_{n_{sp}, n_f}}(sp_1, f_1) \\ S_{11}^{x_{1,1}}(sp_1, f_2) & S_{11}^{x_{1,2}}(sp_1, f_2) & \dots & S_{11}^{x_{2,1}}(sp_1, f_2) & \dots & S_{11}^{x_{n_{sp}, n_f}}(sp_1, f_2) \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ S_{11}^{x_{1,1}}(sp_1, f_{n_f}) & S_{11}^{x_{1,2}}(sp_1, f_{n_f}) & \dots & S_{11}^{x_{2,1}}(sp_1, f_{n_f}) & \dots & S_{11}^{x_{n_{sp}, n_f}}(sp_1, f_{n_f}) \\ S_{11}^{x_{1,1}}(sp_2, f_1) & S_{11}^{x_{1,2}}(sp_2, f_1) & \dots & S_{11}^{x_{2,1}}(sp_2, f_1) & \dots & S_{11}^{x_{n_{sp}, n_f}}(sp_2, f_1) \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ S_{11}^{x_{1,1}}(sp_2, f_{n_f}) & S_{11}^{x_{1,2}}(sp_2, f_{n_f}) & \dots & S_{11}^{x_{2,1}}(sp_2, f_{n_f}) & \dots & S_{11}^{x_{n_{sp}, n_f}}(sp_2, f_{n_f}) \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ S_{11}^{x_{1,1}}(sp_n, f_{n_f}) & S_{11}^{x_{1,2}}(sp_n, f_{n_f}) & \dots & S_{11}^{x_{2,1}}(sp_n, f_{n_f}) & \dots & S_{11}^{x_{n_{sp}, n_f}}(sp_n, f_{n_f}) \end{bmatrix} \quad (3.16)$$

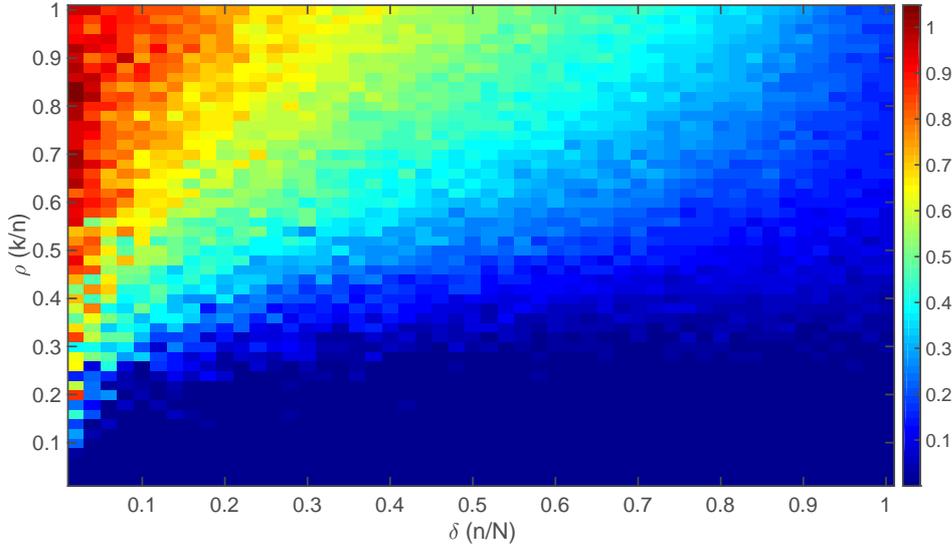
Phase diagram for SAR measurements.

FIGURE 3.11: Phase diagram for SAR measurements with SPGL1. The color scheme depicts the mean relative reconstruction error (blue: reconstruction error is very small, red: high reconstruction error).

In order to verify if CS is a good candidate for the construction of SAR images from incomplete measurement data, we perform now a series of experiments of which the results are summarised in a phase diagram. The phase diagram depicts the mean relative reconstruction error obtained from 100 reconstructions for each combination of the undersampling ratio δ and the sparsity ratio ρ . In order to perform the simulations, we built a measurement matrix A , which models a SAR with a -3 dB antenna opening angle of 30 degrees and emitting SFCW signals over a bandwidth ranging from 1 GHz up to 5 GHz. This synthetic sensor was used to perform the measurement of a scene of 200 pixels. These created scenes were then measured by applying the measurement matrix and reconstructed from the incomplete raw data with the algorithm SPGL-1. Finally, the relative error was obtained by calculating the Euclidean distance between the reconstructed and the original scene.

From the resulting phase diagram 3.11, we conclude that CS can be applied to construct SAR images from incomplete data, if the data is sparse.

3.4 Sequential Compressed Sensing

The compressed Sensing theory states that the number of samples can be reduced far below the lower bound dictated by the Shannon-Nyquist theorem, under certain conditions. This has extensively been proven in numerous applications. Hence, one of the conditions is that the signal to reconstruct must be sparse or has a sparse representation

in a known basis. The mandatory number of samples for a correct reconstruction with a CS reconstruction algorithm is thus, among other things, function of the sparsity or compressibility of the signal to reconstruct. A correct estimation of the sparsity of the signal prior to the measurement is thus necessary to decide on the number of measurements to take. But, in the general case, the signal is unknown prior to the measurement. This makes the estimation of the sparsity difficult and consequently, a correct decision on the number of samples impossible. A solution to this fundamental deadlock situation, is the use of cross-validation measurements and was first proposed in [Ward2009]. The results were obtained through the Johnson-Lindenstrauss lemma [Dasgupta and Gupta2003] for Gaussian and Bernoulli matrices. The work in [Malioutov et al.2010], treats the case of Gaussian measurements and computes explicitly all constants, using a χ^2 distribution. We will now apply this strategy and enlarge this theory to Discrete Fourier Transform (DFT) matrices.

Sequential compressed sensing reconstructs a signal x by sequentially adding samples and performing a new CS reconstruction of x each time, until an imposed reconstruction quality is reached. Compressed Sensing decreases the measurement cost to the expense of an increase in the computation cost. This is even more the case for sequential CS. At each iteration, the signal is constructed until the result passes a cross validation test. For a small increment in number samples and for elevated subsampling rates, many successive reconstructions will be executed, which increases the computation time.

The difficulty for applying this strategy lies in the fact that the ground truth is not available, since we do not have access to x . We can thus not directly compare the reconstructed signal to x , in order to estimate when we are happy with the reconstructed signal. In this section, we will derive a method for calculating an upper bound for the reconstruction error which is sharp enough to be used as stopping criterion for stopping the sequential adding of samples.

3.4.1 Derivation of an upper bound for the reconstruction error

Before we start the derivation and explanation of the proposed technique, we want to put forward that in what follows, we split the complex set of equations $y = Ax$ (where y are the samples resulting from the DFT measurement, A is the DFT measurement matrix and x is the sensed signal), as follows:

$$\begin{bmatrix} Re(y) \\ Im(y) \end{bmatrix} = \begin{bmatrix} Re(A) & -Im(A) \\ Im(A) & Re(A) \end{bmatrix} \begin{bmatrix} Re(x) \\ Im(x) \end{bmatrix}, \quad (3.17)$$

and we introduce the following notation:

$$\mathbf{y} = \mathbf{A}\mathbf{x} \quad (3.18)$$

With

$$\mathbf{y} = \begin{bmatrix} Re(y) \\ Im(y) \end{bmatrix} \quad \mathbf{A} = \begin{bmatrix} Re(A) & -Im(A) \\ Im(A) & Re(A) \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} Re(x) \\ Im(x) \end{bmatrix}.$$

The bold notation denotes thus the split version of the corresponding vector or matrix. The reason for this operation is that the approach that we will elaborate is based on the χ^2 Cumulative Distribution Function (CDF). This distribution function is only defined for real positive numbers and thus, we split the real and imaginary parts of the complex numbers.

The sequential CS reconstruction starts with a first reconstruction $\hat{\mathbf{x}}^{2M}$ obtained from only a small number of M complex samples. We will now derive a strategy to estimate an upper bound for the reconstruction error:

$$\boldsymbol{\delta} = \hat{\mathbf{x}}^{2M} - \mathbf{x}. \quad (3.19)$$

This upper bound will guide us to decide if the reconstruction is successful or if more samples need to be added in order to guarantee a chosen reconstruction quality. The upper bound can be estimated by adding a few T new complex samples, also known as cross-validation measurements, to the set of M samples. The value of the T cross-validation samples can be predicted by performing a synthetic measurement of the reconstructed vector $\hat{\mathbf{x}}^{2M}$, by calculating:

$$\hat{\mathbf{y}}_{2T} = \mathbf{A}_{2T} \hat{\mathbf{x}}^{2M}, \quad (3.20)$$

where \mathbf{A}_{2T} is populated by $2T$ selected rows of \mathbf{A} and which do not belong to the already previously selected $2M$ rows. If the quality of the reconstruction $\hat{\mathbf{x}}^{2M}$ is good enough, we can expect that the predicted measurements $\hat{\mathbf{y}}_{2T}$ are close to the actual measurements \mathbf{y}_{2T} . We will call these differences the deviation errors \mathbf{z}_{2T} :

$$\mathbf{z}_{2T} = \hat{\mathbf{y}}_{2T} - \mathbf{y}_{2T}. \quad (3.21)$$

The deviation errors can be rewritten as:

$$\mathbf{z}_i = \mathbf{A}_i \hat{\mathbf{x}}^{2M} - \mathbf{A}_i \mathbf{x} = \mathbf{A}_i \boldsymbol{\delta}, \quad 1 \leq i \leq 2T, \quad (3.22)$$

with an expectation equal to:

$$\mathbb{E}[\mathbf{z}_i] = \sum_{j=1}^{2N} \mathbb{E}[\tilde{\mathbf{A}}_{ij}] \boldsymbol{\delta}_j \quad (3.23)$$

and a variance:

$$\text{Var}(\mathbf{z}_i) = \boldsymbol{\delta}^T \text{cov}(\tilde{\mathbf{A}}) \boldsymbol{\delta}, \quad (3.24)$$

where $\tilde{\mathbf{A}}$ is obtained by removing the $2M$ previously selected columns from \mathbf{A} . To simplify the problem, we will now replace $\tilde{\mathbf{A}}$ by the full matrix \mathbf{A} (the impact of this decision is discussed further in this section) and we obtain for the expectation:

$$\mathbb{E}[\mathbf{z}_i] = \sum_{j=1}^{2N} \mathbb{E}[\mathbf{A}_{ij}] \boldsymbol{\delta}_j = \frac{1}{2} (\boldsymbol{\delta}_1 + \boldsymbol{\delta}_{N+1}) \quad (3.25)$$

and for the variance:

$$\text{var}(\mathbf{z}_i) = \sum_{j=1}^{2N} \frac{1}{2} \left(1 + \frac{1}{2N-1} \right) \boldsymbol{\delta}_j^2 - \frac{1}{2} \left(1 + \frac{1}{2N-1} \right) \left(\boldsymbol{\delta}_1^2 + \boldsymbol{\delta}_{N+1}^2 \right), \quad (3.26)$$

since the $\text{cov}(\mathbf{A})$ is equal to:

$$\begin{bmatrix} \frac{1}{4} \left(1 + \frac{1}{2N-1}\right) & 0 & \dots & 0 & -\frac{1}{4} \left(1 + \frac{1}{2N-1}\right) & 0 & \dots & 0 \\ 0 & \frac{1}{2} \left(1 + \frac{1}{2N-1}\right) & & & & & & \\ \vdots & & \ddots & & & & & \\ 0 & & & \frac{1}{2} \left(1 + \frac{1}{2N-1}\right) & & & & \\ -\frac{1}{4} \left(1 + \frac{1}{2N-1}\right) & & & & \frac{1}{4} \left(1 + \frac{1}{2N-1}\right) & & & \vdots \\ 0 & & & & & \frac{1}{2} \left(1 + \frac{1}{2N-1}\right) & & \\ \vdots & & & & & & \ddots & 0 \\ 0 & & & & & & & \frac{1}{2} \left(1 + \frac{1}{2N-1}\right) \end{bmatrix} \quad (3.27)$$

with:

- $\frac{1}{4} \left(1 + \frac{1}{2N-1}\right)$ on positions $(1,1)$ and $(N+1, N+1)$.
- $-\frac{1}{4} \left(1 + \frac{1}{2N-1}\right)$ on positions $(N+1,1)$ and $(1, N+1)$.

We further suppose that we can make following approximations:

1. $E[\mathbf{z}_i] = 0$
2. $\text{var}(\mathbf{z}_i) = \frac{1}{2} \|\boldsymbol{\delta}\|_2^2$, if N is large.

We then know that (thanks to the central limit theorem):

$$\frac{\sum_{i=1}^{2T} \mathbf{z}_i^2}{\frac{1}{2} \|\boldsymbol{\delta}\|_2^2} \sim \chi_{2T}^2. \quad (3.28)$$

From the χ_{2T}^2 Cumulative Distribution Function (CDF) we can obtain the largest z^* , for a small α ($0 \leq \alpha \leq 1$), such that:

$$P\left(\frac{\sum_{i=1}^{2T} \mathbf{z}_i^2}{\frac{1}{2} \|\boldsymbol{\delta}\|_2^2} \leq z^*\right) \leq \alpha. \quad (3.29)$$

This equation can be reorganised to obtain an upper bound for $\|\boldsymbol{\delta}\|_2$:

$$P\left(\|\boldsymbol{\delta}\|_2 \leq \sqrt{\frac{2 \sum_{i=1}^{2T} \mathbf{z}_i^2}{z^*}}\right) \geq 1 - \alpha. \quad (3.30)$$

In other words: we can deduce, with a probability equal to $1 - \alpha$, an upper bound for the ℓ_2 -norm of the reconstruction error from a small number of T cross-validation measurements without having access to the ground truth x .

3.4.2 Influence of the approximations

During the deduction of the upper bound, we introduced several approximations. Without approximations, the expression in Equation (3.28) becomes:

$$\frac{\sum_{i=1}^{2T} (\mathbf{z}_i - \sum_{j=1}^{2N} \mathbb{E}[\tilde{\mathbf{A}}_{ij}] \boldsymbol{\delta}_j)^2}{\boldsymbol{\delta}^T \text{cov}(\tilde{\mathbf{A}}) \boldsymbol{\delta}} \sim \chi_{2T}^2. \quad (3.31)$$

We will now empirically test the impact of using \mathbf{A} instead of $\tilde{\mathbf{A}}$ and of the approximations $\mathbb{E}[\mathbf{z}_i] = 0$ and $\text{var}[\mathbf{z}_i] = \frac{1}{2} \|\boldsymbol{\delta}\|_2^2$, by verifying if the expression in (3.28), follows a χ_{2T}^2 distribution. This verification is performed by applying a Kolmogorov-Smirnov (KS) test [[Lilliefors1967]] which returns a test decision for the null hypothesis that the data comes from a χ_{2T}^2 distribution. We define the outcome equal to 1 if the test rejects the null hypothesis at the 5% significance level, or 0 otherwise. This test is performed for all combinations of the signal length N , the number of samples used for the reconstruction n and the sparsity k of the original signal. The phase diagrams are thus an appropriate choice for representing the results of these tests.

The empiric tests are conducted as follows:

1. We create a vector x of length $N = 512$ and with a sparsity k . The positions of the k non-zero elements are randomly picked and their values come from a standard normal distribution.
2. We perform a CS reconstruction of x , using SPGL-1, from $n = 2M$ randomly selected samples.
3. We then randomly select 100 cross-validation samples out of the ensemble of samples which were not yet used for the CS-reconstruction and calculate the corresponding deviation errors \mathbf{z}_i .
4. From the deviation errors, we deduce the Probability Density Functions (PDFs) and the CDFs of:

(a)

$$\frac{(\mathbf{z}_i - \sum_{j=1}^{2N} \mathbb{E}[\tilde{\mathbf{A}}_{ij}] \boldsymbol{\delta}_j)^2}{\boldsymbol{\delta}^T \text{cov}(\tilde{\mathbf{A}}) \boldsymbol{\delta}} \quad (3.32)$$

(b)

$$\frac{\mathbf{z}_i^2}{\frac{1}{2} \|\boldsymbol{\delta}\|_2^2}. \quad (3.33)$$

5. The calculated CDFs are submitted to the previously described Kolmogorov-Smirnov test, which has as outcome the test-decision and the p -value.
6. This test is repeated 30 times for all combinations of the subsampling ratio $\delta (= n/N)$ and the sparsity ratio $\rho (= k/n)$.

The outcome of these tests are depicted in Figure 3.12:

- In (a) we show the test decision of the KS-test on the CDF obtained from (3.32). We observe that the null-hypothesis is accepted in the large majority of combinations

of ρ and δ as long as δ is not too large (approximately $> 80\%$). If δ is too large, the selected cross-validation samples are no longer independent since they are selected from a fairly small ensemble and consequently (3.32) does no longer pass the KS-test. We can conclude that (3.32) follows a χ^2_2 distribution if δ is not too large (which is, in general, the case for CS).

- This result is confirmed in **(b)**, where we depict the p -value over 30 independent tests, with each time a new random vector x and this for all combinations of δ and ρ .
- In **(c)**, we depict the outcome of the KS-tests on the CDFs obtained from (3.33). We see that the expression in (3.33) does not follow a χ^2_{2T} as soon as δ is larger than approximately 20%.
- This result is again confirmed in **(d)** where we show the corresponding mean p -values over 30 independent executions of the test. As soon as δ increases, the p -value drops.

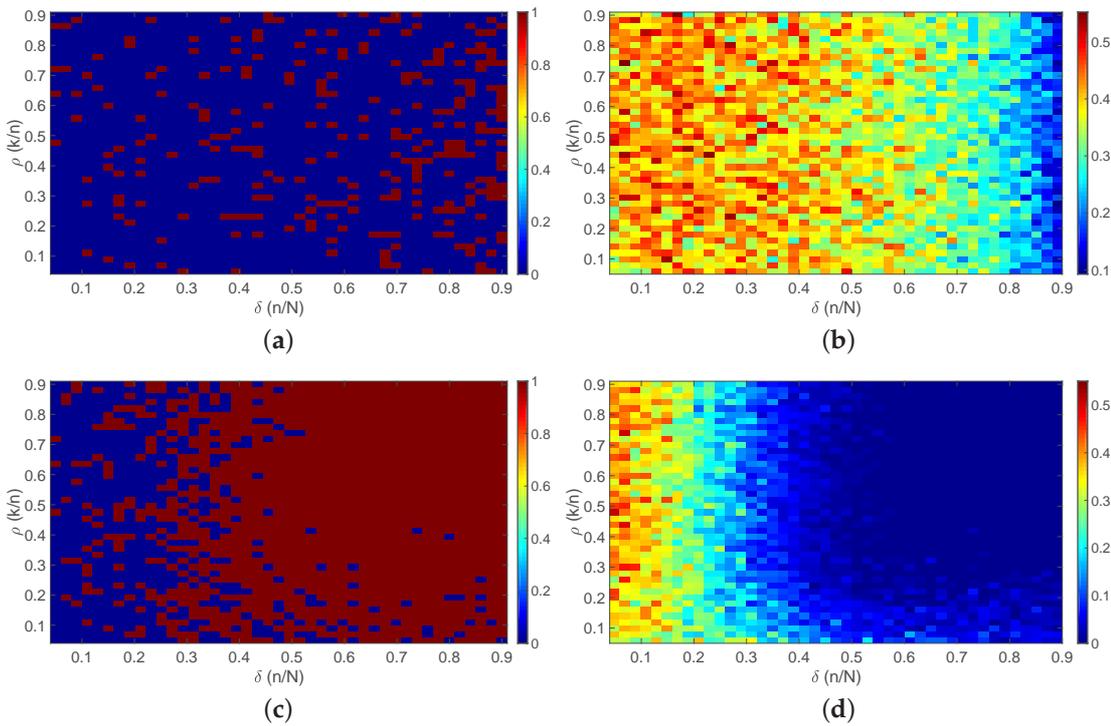


FIGURE 3.12: Phase diagrams depicting the outcome of the KS-test on the null-hypothesis that the expressions in (3.32) and (3.33) are drawn from a χ^2_{2T} distribution. In **(a)** and **(b)** are respectively the test decisions (1: rejects the null-hypothesis and 0: accept the null-hypothesis) and the mean p -values (over 30 executions) for (3.32). In **(c)** and **(d)** are the test decisions and the p -values for (3.33).

From these tests, we conclude that the expression in (3.28) does not hold for increasing values of δ . This is because we supposed that we pick the cross-validation samples out of

the total ensemble of samples. Whereas, in reality, we chose the cross-validation samples out of the $2N - 2M$ samples which were not selected to construct the signal. The more samples used for the reconstruction (and thus, the larger is δ), the more important becomes the difference between $\tilde{\mathbf{A}}$ and \mathbf{A} and thus the higher the impact of the introduced error. We could thus conclude that the method is only applicable for fairly small subsampling rates. But, before making that conclusion, let us investigate what happens to the CDF of (3.33) when the subsampling rate increases.

The impact of an increasing subsampling rate on the CDF of the expression in (3.28) is illustrated through the following test:

1. First, we create a vector x of length $N = 256$. We choose a subsampling ratio $\delta = 5\%$ and a sparsity ratio $\rho = 20\%$. The positions of the non-zero elements in x are picked randomly and their values are drawn from a standard normal distribution.
2. The original vector x is reconstructed, using the SPGL-1 CS-reconstruction algorithm, from n samples, randomly selected from a DFT-measurement of the vector x .
3. Subsequently, we randomly select $T = 8$ cross-validation samples out of the ensemble of samples which were not used for the reconstruction. This selection is repeated independently 100 times.
4. The resulting 100 deviation errors \mathbf{z}_{2T} are then used to estimate the CDFs obtained from:
 - (a) $\frac{\sum_{i=1}^8 (\mathbf{z}_i - \mathbb{E}[\mathbf{z}_i])^2}{\delta^T \text{cov}(\tilde{\mathbf{A}}) \delta}$ (using $\tilde{\mathbf{A}}$ and without approximations),
 - (b) $\frac{\sum_{i=1}^8 (\mathbf{z}_i - \mathbb{E}[\mathbf{z}_i])^2}{\delta^T \text{cov}(\mathbf{A}) \delta}$ (using \mathbf{A} and without approximations),
 - (c) $\frac{\sum_{i=1}^8 \mathbf{z}_i^2}{\frac{1}{2} \|\delta\|_2^2}$ (using \mathbf{A} and with approximations)
5. This test is repeated with subsampling ratios $\delta = 25\%$ and $\delta = 65\%$.

The CDFs resulting from this experiment, are depicted in Figure (3.13). The blue curve corresponds to the theoretical χ_{2T}^2 distribution for $T = 4$. The other CDF-curves are obtained from the measured deviation errors. The red curve is obtained when using the matrix $\tilde{\mathbf{A}}$ without approximating the expectation and variance of \mathbf{z}_i . The yellow curves show the CDFs when replacing $\tilde{\mathbf{A}}$ by \mathbf{A} (without applying the approximations). Whereas the magenta markers correspond to the CDFs obtained with \mathbf{A} and the approximations on $\mathbb{E}[\mathbf{z}_i]$ and $\text{var}(\mathbf{z}_i)$.

On the graphs in Figure (3.13) we observe:

- The approximations do not have a relevant impact on the PDF or the CDF.
- Secondly, we see the following impact of increasing the subsampling rate when replacing $\tilde{\mathbf{A}}$ with \mathbf{A} : The CDFs obtained with \mathbf{A} are upper bounded by the CDF obtained with $\tilde{\mathbf{A}}$. For increasing subsampling rates the PDFs are shift to the right. This means that the CDF will decrease for the same value when the subsampling rate increases if \mathbf{A} is used.

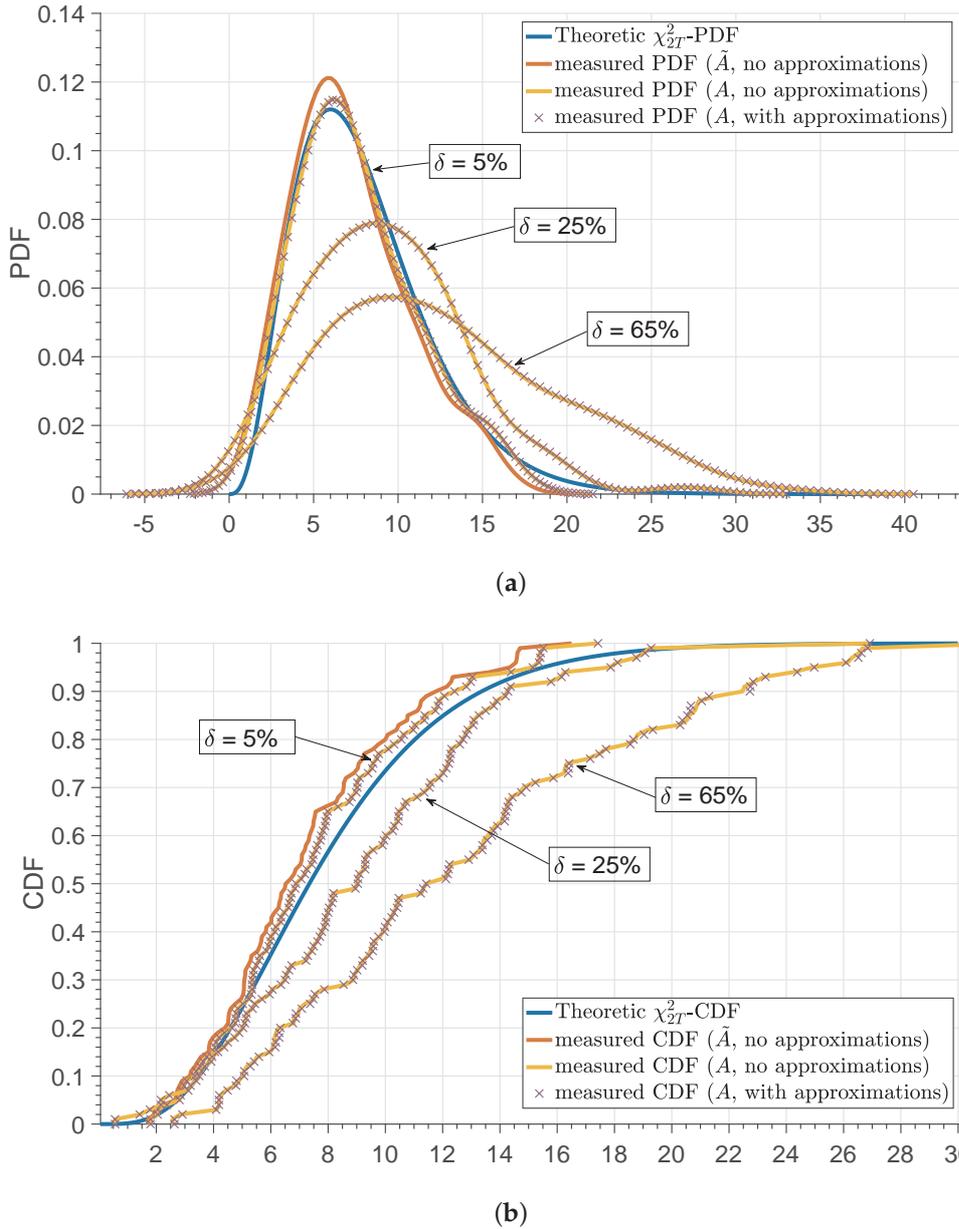


FIGURE 3.13: In (a) the PDFs and in (b) the corresponding CDFs, with in blue: the PDF/CDF of the theoretical χ^2_{2T} distribution with $T = 4$. In red: the PDF/CDF obtained from the expression in (3.32). In yellow: the PDF/CDF obtained when replacing \tilde{A} with A in (3.32). In magenta: the PDF/CDF obtained when replacing \tilde{A} with A in (3.32) and with the approximations on $E[z_i]$ and $\text{var}(z_i)$.

The impact of the approximations can thus be neglected, whereas using A instead of \tilde{A}

will have the following effect on the estimated upper bound for the reconstruction error: Because the CDF is lower than the theoretical CDF (which is used to determine the value z^*), the value of z^* in the expression of the upper bound (Equation 3.30) will be too small and thus, the upper bound will increase. In other words, the upper bound obtained in 3.30 is an overestimation of the upper bound for the reconstruction error. We will show that the upper bound is still sharp enough to be relevant.

3.4.3 An example

We will now empirically evaluate the method for estimating the upper bound as a whole. First, by way of example, we test the method for the reconstruction of two signals with a different sparsity ratio, as follows:

1. We create two vectors x_1 and x_2 of length $N = 2000$ and $k_1 = 60$ and $k_2 = 120$ non-zero elements respectively. The positions of the non-zero elements in x_1 and x_2 are randomly chosen and their values are drawn from a standard normal distribution.
2. The vectors x_1 and x_2 are reconstructed, using the SPGL-1 CS reconstruction algorithm, from an increasing number of samples which are randomly selected from the DFT-measurements of x_1 and x_2 .
3. We then calculate the reconstruction errors $\|\delta\|_2$ ($k = 60$) and $\|\delta\|_2$ ($k = 120$) and determine the upper bound for the reconstruction errors (using Equation (3.30)) with $T = 20$ randomly selected cross-validation samples and $\alpha = 0.05$.

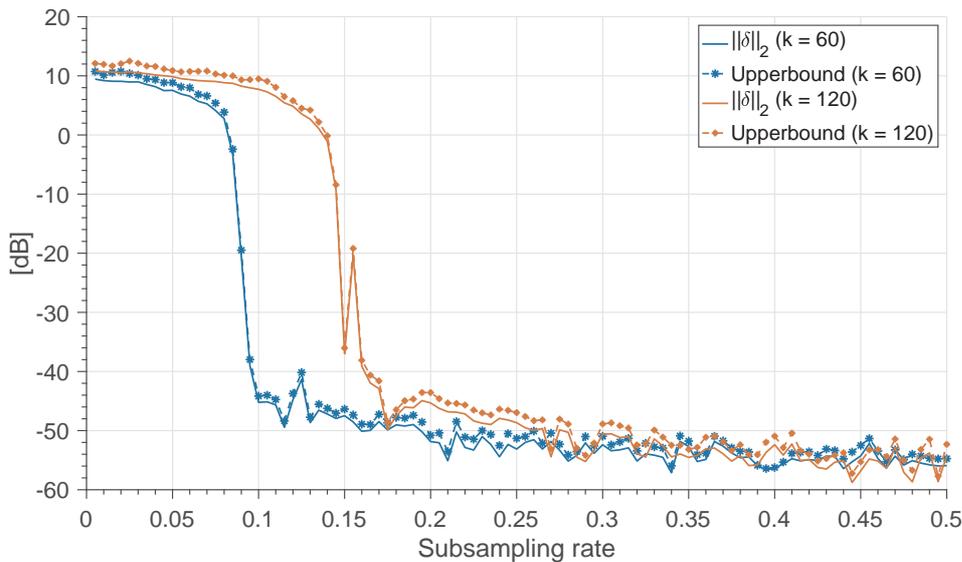


FIGURE 3.14: Reconstruction error for increasing sampling rates (in blue: $k = 60$ and in red: $k = 120$). The corresponding upper bounds are represented by the dashed curves.

The results of this experiment are depicted in Figure 3.14. We observe:

1. The reconstruction error is correctly and neatly bounded by the estimated upper bound.
2. Although the upper bound is selected with a probability equal to 95%, the upper bound never drops below the true error. This can be explained by the fact that the upper bound is overestimating the error because we use \mathbf{A} instead of $\tilde{\mathbf{A}}$.

3.4.4 Empiric evaluation of the proposed method

We now repeat the same experiment for all possible combinations of the subsampling ratio $\delta (= n/N)$ and the sparsity ratio $\rho (= k/n)$. The results are depicted in the phase diagrams in Figure 3.15. In **(a)**, we see the true reconstruction errors and in **(b)** we have the corresponding estimated upper bounds for the reconstruction error ($T = 10$ and $\alpha = 0.05$). In **(c)** we represent the difference between the reconstruction error and the upper bounds. We only observe a significant difference in the region where the reconstruction errors are the highest. In **(d)**, finally, the combinations of the subsampling ratio and the sparsity ratio where the upperbound is lower than the true error, are indicated in red. We count only 42 (which is much lower than 5% of the in total 2500 combinations) upper bounds which happen to be smaller than the corresponding true reconstruction error.

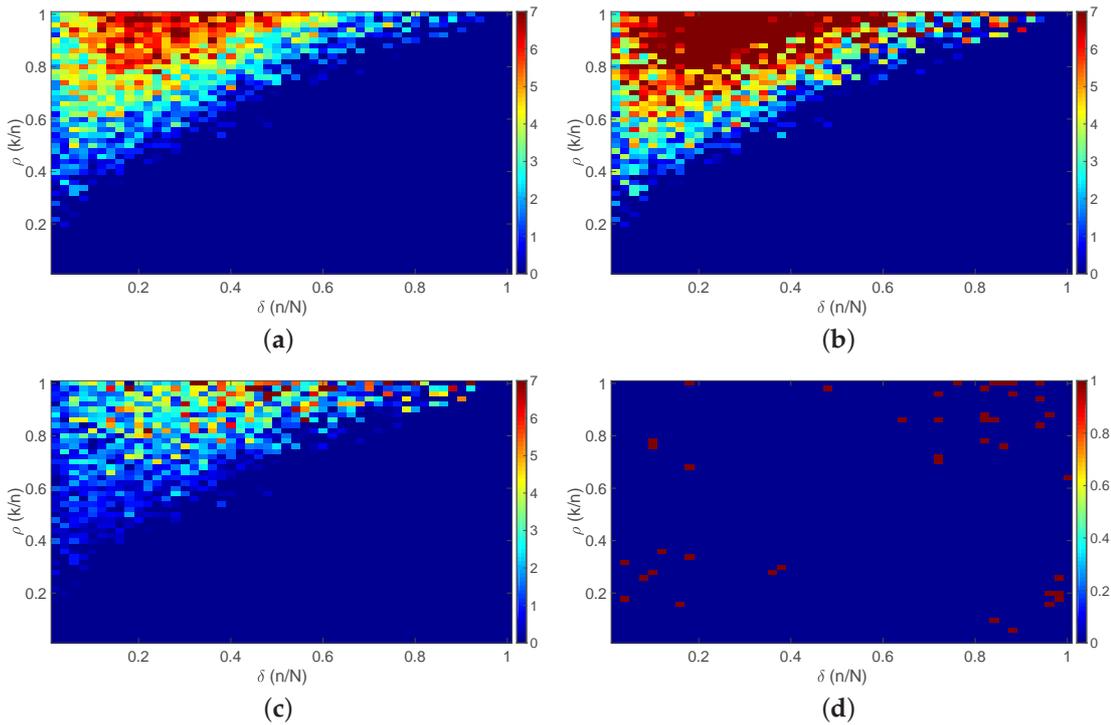


FIGURE 3.15: **(a)** Phase diagram with the true reconstruction error. **(b)** Estimated upper bound for the reconstruction error (using $T = 10$ cross-validation measurements and $\alpha = 0.05$). **(c)** Difference between the true reconstruction error and the estimated upper bounds. **(d)** In red: combinations of ρ and δ where the upper bound is lower than the true error.

Finally, we repeat the same test 30 times. The phase diagram in Figure 3.16 **(a)** shows the mean reconstruction error for all combinations of δ and ρ over the 30 iterations. In **(b)**, we represent the phase diagram with the mean of the upper bound over the 30 iterations. In **(c)** we depict the mean of the difference between the reconstruction error and the upper bound. In **(d)**, finally, we depict the variance of the difference between the reconstruction error and the upper bound.

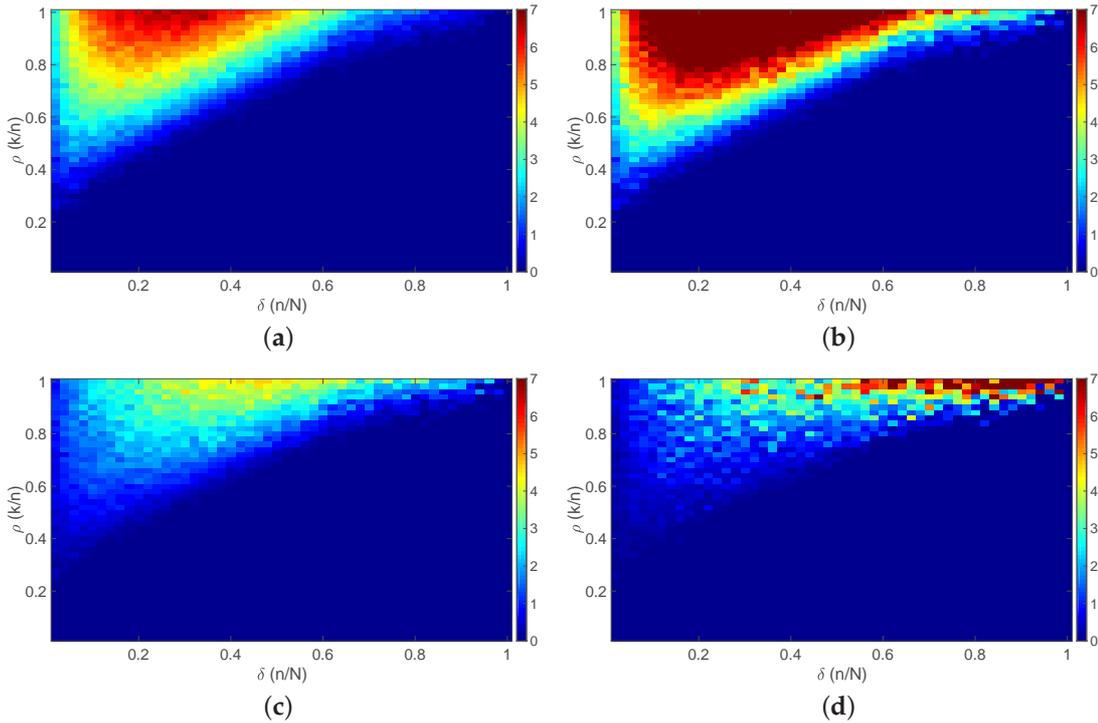


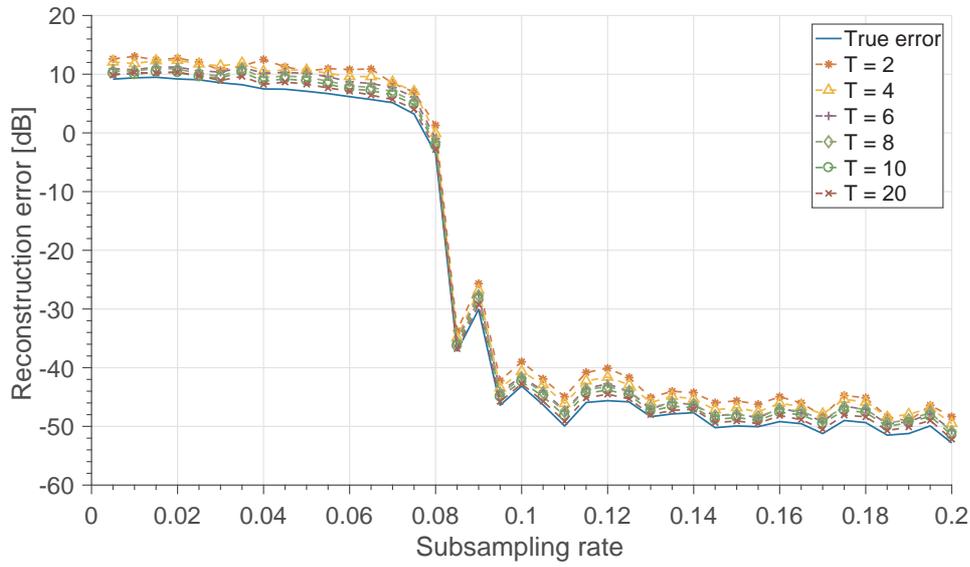
FIGURE 3.16: **(a)** Phase diagram with the mean reconstruction error for all combinations of ρ and δ over 30 iterations. **(b)** Phase diagram with the mean of the upper bounds over the 30 iterations. **(c)** Phase diagram with the mean of the difference between the true error and the upper bound. **(d)** Phase diagram with the variance of the difference between the true error and the upper bound.

All of the experiments on synthetic data, up till now, illustrate that the proposed method allows to estimate an usable upper bound for the reconstruction error.

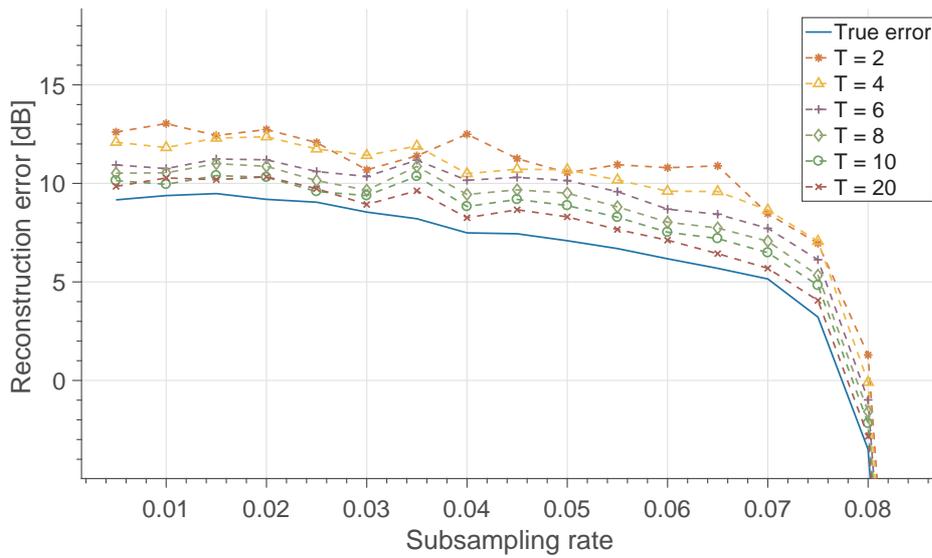
3.4.5 Number of cross-validation measurements and the upper bound probability

We will now evaluate the impact of the number of cross-validation measurements and the choice of the confidence parameter α on the upper bound for the reconstruction error $\|\delta\|_2$. For these experiments, we neglect the fact that we use the $\text{var}(\mathbf{A})$ instead of $\text{var}(\tilde{\mathbf{A}})$ to estimate the upper bound. We know that the upper bound is overestimated by doing so. We perform again a set of synthetic measurements of a vector of length $N = 200$ and with $k = 30$ nonzero elements. The reconstruction at the first iteration of the sequential CS approach is obtained from only 6 randomly picked samples. At each subsequent iteration we add 6 new samples to the set of already selected samples. At each iteration we plot (Figure 3.17): (1) the actual measured error (the solid blue curve), and (2) the calculated error bound (dashed curves), obtained from a different number of cross-validation measurements $T = 2, 3, 6, 8, 10$ and 20 with $1-\alpha = 0.9$. We can conclude from the obtained results that, as can be expected, the estimated error bounds become sharper when the

number of cross-validation samples increases. Nevertheless, we already obtain a useful upper bound from a very small number of cross-validation samples.



(a)

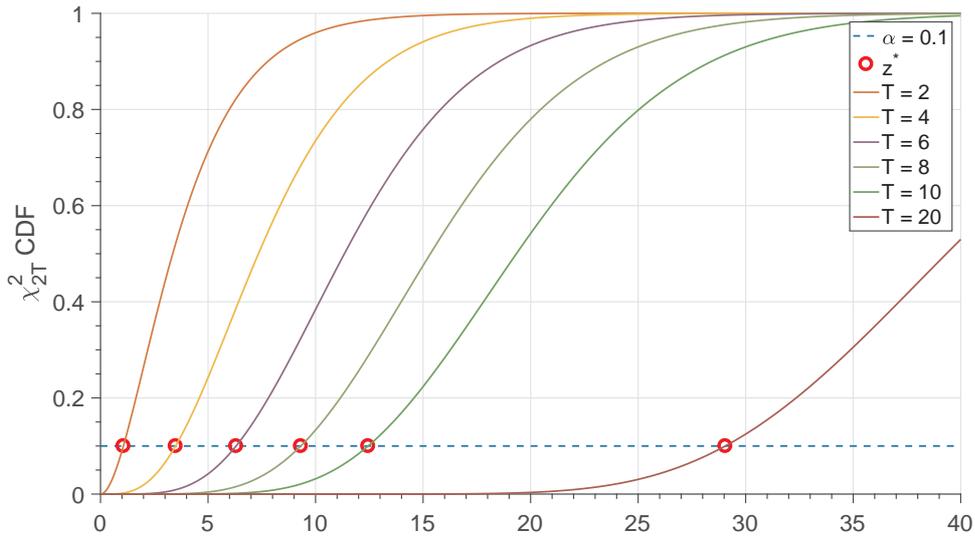


(b)

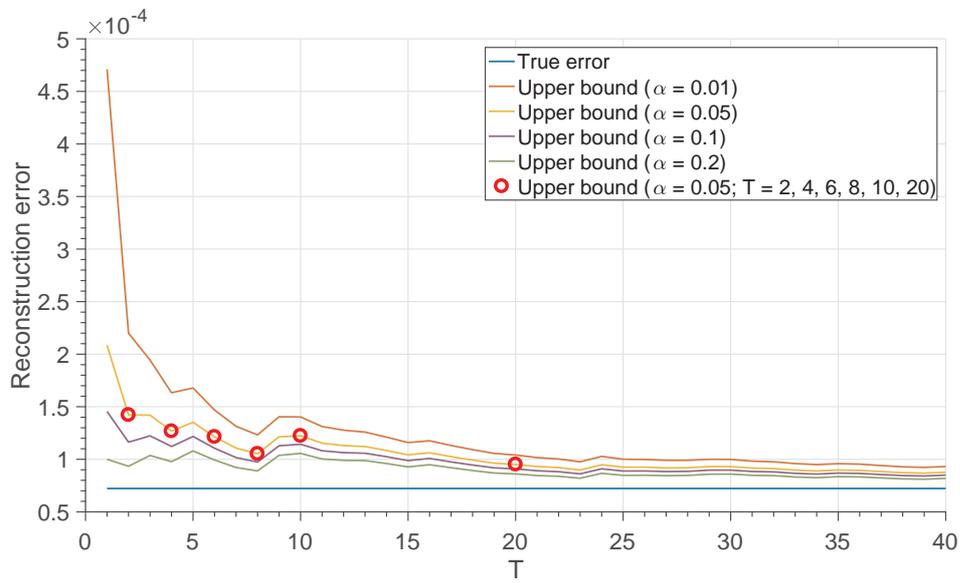
FIGURE 3.17: (a) The measured true error δ (solid blue curve) and the estimated bounds for an increasing number of measurements using a different number of cross validation measurements (T). (b) Zoom of figure (a).

In Figure 3.18 (a) we depict the χ_{2T}^2 CDFs for increasing numbers of independent cross-validation measurements T , which corresponds to increasing the degrees of freedom of the χ^2 distribution. We see, indeed, that for larger values of T , z^* is obtained at larger values of z for the same chosen value for α . Consequently, the estimated bound $\sqrt{\frac{2\sum_{i=1}^T z(i)^2}{z^*}}$ decreases and becomes thus a sharper upper bound for $\|\delta\|_2$. The obtained z^* for $\alpha = 0.1$ (which were used to calculate the upper bounds depicted in Figure 3.17) are encircled in red.

Figure 3.18 (b) shows the reconstruction error and the estimated error bounds from a fixed number of samples equal to 80, for increasing numbers of cross-validation samples and for different values of α . We observe that the error bounds rapidly decrease towards the actual error. We further observe that the error bounds decrease for increasing values of α . We can conclude from these experiments that the number of cross-validation measurements can be kept below 5% of the number of the signal length, even for low values of α .



(a)



(b)

FIGURE 3.18: (a) CDF's of the χ^2_T distribution function for increasing T and the z^* for $\alpha = 0.1$ are encircled in red. (b) True measured error, the estimated error and the estimated error bounds obtained with the z^* from (a) with $\alpha = 0.1, 0.01, 0.05$ and 0.2 .

3.5 Preconditioning methods

3.5.1 Singular Value Decomposition

The Singular Value Decomposition (SVD) of a real or complex matrix $M \in \mathbb{C}^{m \times n}$ corresponds to the following factorization:

$$M = U\Sigma V^*, \quad (3.34)$$

with $U \in \mathbb{C}^{m \times m}$ and $V \in \mathbb{C}^{n \times n}$ unitary matrices whose columns are the left-singular vectors and right-singular vectors respectively. $\Sigma \in \mathbb{C}^{m \times n}$ is a rectangular diagonal matrix with the singular values of M on the diagonal entries σ_i . V^* denotes the conjugate transpose of V .

In the frame of this work, the SVD of a matrix can be applied in three distinct ways:

(1) *SVD used as a data-adaptive sparsity basis:*

The work in [Hong et al.2011] proposes to use the SVD as a sparsifying basis for MRI imaging. The proposed method works as follows: From the subsampled measurements a first approximation (I_1) is calculated through the inverse Fourier transform. This operation results in an aliased version of the object image. Subsequently an SVD of I_1 (3.34), with the singular values on the diagonal of Σ_1 arranged in non-increasing order: $\text{diag}(\sigma_{11}, \sigma_{12}, \dots, \sigma_{1r}, \dots, 0)$, can be approximated by abandoning the smallest singular values (u_{1i} and v_{1i} are the corresponding vectors of U and V):

$$I_1 \approx I_1^k = \sum_{i=1}^k \sigma_{1i} u_{1i} v_{1i}^T, \quad k \leq r. \quad (3.35)$$

The corresponding matrices U_1 and V_1 are the initial estimates for the sparsifying transform [Hong et al.2011]:

$$\Psi_1(x) = U_1 x V_1^* \quad (3.36)$$

A new and better approximation of the object MRI image I_2 is then obtained using a CS reconstruction algorithm together with the sparsifying transform (3.36). An updated version of U and V can be deduced from the reconstructed Image I_2 which allows to calculate again a better estimation of I and so on. Due to the adaptive character of the data, this approach will sparsify a broader range of images than predefined transforms.

(2) *Truncated SVD of the measurement matrix:*

By truncating the SVD decomposition of the measurement matrix A , a much slimmer measurement matrix can be obtained and used together with a CS algorithm to reconstruct the original signal. Applying a SVD in this way provides advantages in terms of computational time while maintaining the quality of the full scene reconstruction. The critical step in this procedure is the truncation or the choice of the number of singular values to keep. The idea was proposed and tested in [Zhang et al.2015] for a down-looking linear array 3-D SAR.

(3) *Clutter suppression:*

In common signal processing applications, the SVD decomposition is used to obtain a

compressed signal where the dominant singular values are kept and the smallest singular values are considered to be noise and thus neglected. For some applications however the opposite is true: the information of interest is hidden in the smaller eigenvalues and is masked by a stronger signal. Clutter can be a performance limiting factor in radar applications as is the case for Ground-Penetrating Radars (GPR) [Gunatilaka and Baertlein2000] applied for the detection of near-surface targets including anti-personnel mines and Through-the-Wall Radar Imaging (TWI) [Tivive et al.2014, Tivive et al.2011b, Bouzerdoum and Tivive2015]. For both applications, the strong clutter coming from the surface or the wall is supposed to be contained within the first eigenvalue(s) of the SVD of the raw data. By eliminating these eigenvalues from the SVD, the strong clutter signal is suppressed.

3.5.2 Principal Component Analysis

Principal Component Analysis (PCA) is a widely used statistical tool for data pattern detection and dimensionality reduction. It is an orthogonal transformation, deduced from a representative training set, to transform a set of samples of correlated variables into a set of linear uncorrelated variables which are the principal components. Suppose the matrix $M \in \mathbb{C}^{N \times n}$ which columns are populated by n training vectors. After subtracting of the mean value from each training vector, we calculate the covariance matrix of M . The eigenvectors of the covariance matrix provide us with information about the patterns in the data and are used to populate the orthogonal transformation matrix. The significance of the components is determined by the corresponding eigenvalues. The PCA transformation matrix can thus be limited to those components corresponding to the largest eigenvectors.

(1) *PCA used as a sparsity basis:*

PCA can be used as a sparsifying transform for the CS reconstruction of SAR images [Wang et al.2012, Pourkamali-Anaraki and Becker2017]. This approach was recently tested and evaluated on real SAR Ground Penetrating Radar (GPR) data in [Cristofani et al.2018]. The technique can be summarized as follows:

The minimization problem (2.10) is converted to:

$$\min_x \|x\|_1 \quad \text{s.t.} \quad \|\tilde{A}x - \tilde{y}\|_2 \leq \epsilon, \quad (3.37)$$

where $\tilde{A} = T_A^* A$ and $\tilde{y} = T_A^* y$. The transformation matrix T_A is obtained through the PCA of the training set of possible returns which is described by the matrix A and is limited to the largest corresponding eigenvalues.

(2) *Robust Principal Component Analysis:*

Robust Principal Component Analysis (RPCA) is a different approach for integrating PCA into the CS reconstruction problem. The method was distinctly proposed by Candès et al. in [Candès et al.2011] and by Chandrasekaran et al. in [Chandrasekaran et al.2011] and solves the following separation problem: Suppose that the matrix M can be decomposed as:

$$M = L_0 + S_0, \quad (3.38)$$

with L_0 a low rank matrix and S_0 a sparse matrix, can we recover both low-rank and sparse components from an incomplete set of samples? The work in [Candès et al.2011] proves that this can be done under rather weak assumptions, by solving the following Principal Component Pursuit problem:

$$\min_{L,S} \|L\|_* + \lambda \|S\|_1 \quad \text{s.t.} \quad L + S = M. \quad (3.39)$$

With λ a regularization parameter and where $\|L\|_*$ denotes the nuclear norm of L :

$$\|L\|_* := \sum_i \sigma_i(L), \quad (3.40)$$

with σ_i the singular values of L . The work in [Candès et al.2011] demonstrates that $1/\sqrt{n}$ is a good choice for the regularization parameter λ . This method allows thus to recover the low-rank matrix L_0 from an undersampled set of measurements corrupted by S_0 which happens to be sparse. The technique has been applied for the detection of objects in a cluttered background for video surveillance, for face recognition [Candès et al.2011], for moving target detection [Cao et al.2016] and for radar Moving Target Indication [Yang et al.2015].

3.6 Compressed Sensing with Side Information

3.6.1 Homogeneous Side Information

The Shannon-Nyquist sampling theorem fixes the lower bound for the number of samples needed for a correct reconstruction of a signal. This bound can only further be lowered if additional information on the signal is known prior to the measurement. The cornerstone of Compressed Sensing is the prior supposition that the signal will be sparse or has a sparse representation in a known basis. As seen before, this knowledge can help to drastically decrease the sampling rate. In many applications, sparseness is not the only prior knowledge available. In this section we explore how to add prior knowledge under the form of a similar signal, also known as Side Information (SI), in the reconstruction of a sparse signal. We will test and compare three strategies to add the SI:

- Coherent background subtraction [Cevher et al.2008] is a well-known and effective way to insert the SI into the reconstruction algorithm. This approach was successfully used for example for the reconstruction of MR images [Trzasko et al.2009] and in through-the-wall radar imaging [Martone et al.2010]. Consider the vector $x \in \mathbb{C}^N$ to be the unknown vector we want to reconstruct from a subsampled measurement $y \in \mathbb{C}^N$. Suppose further that $z \in \mathbb{C}^N$ is a signal with a high degree of similarity compared to x and which was antecedently obtained from a full measurement Az . The coherent background subtraction technique reconstructs the difference between the signals x and z , which can be supposed to be sparser than the signal x itself, if the signals are highly similar. Coherent background subtraction neglects thus the sparsity of x and concentrates on minimizing the number of non-zero elements of the foreground ($x - z$). The reconstruction can be obtained after coherently subtracting the measurements Az from the samples of the measurement y . The solution is obtained by solving the following minimization problem:

$$\min \|x - z\|_1 \quad \text{s.t.} \quad \|A(x - z) - (y - Az)\|_2 \leq \epsilon \quad (3.41)$$

- $\ell_1\ell_n$ -Minimization is a different approach for adding the SI z into the minimization equation and integrates the prior information alongside the sparsity assumption into the CS algorithm [Mota et al.2017, Mota et al.2015]. Here, we want thus to minimize the weighted sum of the ℓ_1 -norm of x and the difference between x and z :

$$\min\{\|x\|_1 + \beta g(x - z)\} \quad \text{s.t.} \quad \|Ax - y\|_2 \leq \epsilon, \quad (3.42)$$

where $g(\cdot)$ is a function measuring the similarity between x and z and $\beta (> 0)$ is a trade-off factor between the sparsity of x and the similarity between x and z . In [Mota et al.2014], two common choices for $g(\cdot)$ are compared and evaluated: the ℓ_2 -norm (resulting in $\ell_1\ell_2$ -minimization) and the ℓ_1 -norm (resulting in $\ell_1\ell_1$ -minimization). The authors prove that the $\ell_1\ell_1$ approach outperforms the $\ell_1\ell_2$ -minimization and that the sharpest undersampling bound is obtained for $\beta = 1$:

$$\min\{\|x\|_1 + \|x - z\|_1\} \quad \text{s.t.} \quad \|Ax - y\|_2 \leq \epsilon \quad (3.43)$$

- The $\ell_1\ell_1$ -minimization technique can easily be extrapolated towards a CS reconstruction problem with multiple SI. The minimization problem (3.43) becomes:

$$\min_x \left\{ \frac{1}{2} \|Ax - y\|_2^2 + \lambda \sum_{j=0}^J \|x - z_j\|_1 \right\}, \quad (3.44)$$

where J SIs are added to the CS minimization problem and where $z_0 = 0$ and where λ denotes the Lagrange multiplier [Bertsekas2014]. This simple approach attributes the same importance to each of the SIs, regardless of the difference in similarity or quality. SI of poor quality will thereby corrupt a correct reconstruction of the signal x . Alternatively, the work in [Van Luong et al.2016] presents an approach which attributes weights at two different levels to the SIs: (1) inter-SI weights $\beta_j \geq 0$ between the different SIs and (2) intra-SI weights w_{ji} , among the different elements inside a single SI. Equation (3.44), becomes:

$$\min_x \left\{ \frac{1}{2} \|Ax - y\|_2^2 + \lambda \sum_{j=0}^J \beta_j \|W_j(x - z_j)\|_1 \right\}, \quad (3.45)$$

where $W_j \geq 0$ is a diagonal matrix belonging to the SI z_j and whose diagonal elements ($w_{j1}, w_{j2}, \dots, w_{jn}$) are populated with the intra-SI weights of z_j . When choosing $\beta_j = 1$ and $W_j = I$ (where I is the unit matrix), the minimization problem (3.45) is reduced to the non-weighted problem (3.44). The weights β_j will valorize the global impact of the SIs processing a high degree of similarity compared to x and suppress the impact of poor quality SI. The weights W_{ji} , on the other hand, regulate the relative impact of the different elements inside a single SI z_j . In [Van Luong et al.2016], a Reconstruction Algorithm with Multiple SI using Adaptive weights (RAMSIA) that solves the minimization problem (3.45), is presented. RAMSIA is based on the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [Beck and Teboulle2009], to which the iterative computation of the weights is added. Each iteration contains two steps for obtaining the weights: (1) W_j is computed, given x and β_j , by optimizing:

$$\underset{w_{ji}}{\operatorname{argmin}} \left\{ \lambda \beta_j \sum_{i=0}^n w_{ji} |x_i - z_{ji}| \right\}, \quad (3.46)$$

where z_{ji} is an element of z_j at index i . (2) Considering x and W_j fixed, we compute β_j , by calculating:

$$\operatorname{argmin}_{\beta_j} \left\{ \lambda \sum_{j=0}^J \beta_j \|W_j(x - z_j)\|_{\ell_1} \right\}. \quad (3.47)$$

Finally, we compute $x^{(k)}$, the solution for (3.45) at iteration k given the obtained W_j and β_j . We refer the reader to [Van Luong et al.2016] for the pseudo code of the RAMSIA algorithm.

We will now test and evaluate the performance of the common CS reconstruction without SI, the coherent background subtraction technique and the approach combining CS with the weighted SI (RAMSIA) on simulated data. The synthetic experiments are performed using a SAR measurement matrix with similar parameters as the sensor that will be used for the Non-Destructive Tests (NDT) in chapter 4. This measurement matrix is further used to reconstruct the sensed random k -sparse ($k = 10$) vectors of length $N = 200$. The reconstruction quality is measured as the Euclidean distance between the original vector x and its reconstructed version, for a sampling rate starting at 2% and increasing up to 60% (with a step of 2%). To avoid dependency of the selected samples, the tests are executed 32 times using a new pseudo-random sampling scheme for each execution of the test. The depicted results are the mean reconstruction errors over the 32 executions of the test together with the 95% symmetric confidence intervals. (The confidence interval of n instances of variable x is thus: $[\bar{x} - t^* \frac{\sigma}{\sqrt{n}}, \bar{x} + t^* \frac{\sigma}{\sqrt{n}}]$, with \bar{x} the mean, σ the standard deviation and t^* is obtained from the Student's t distribution.)

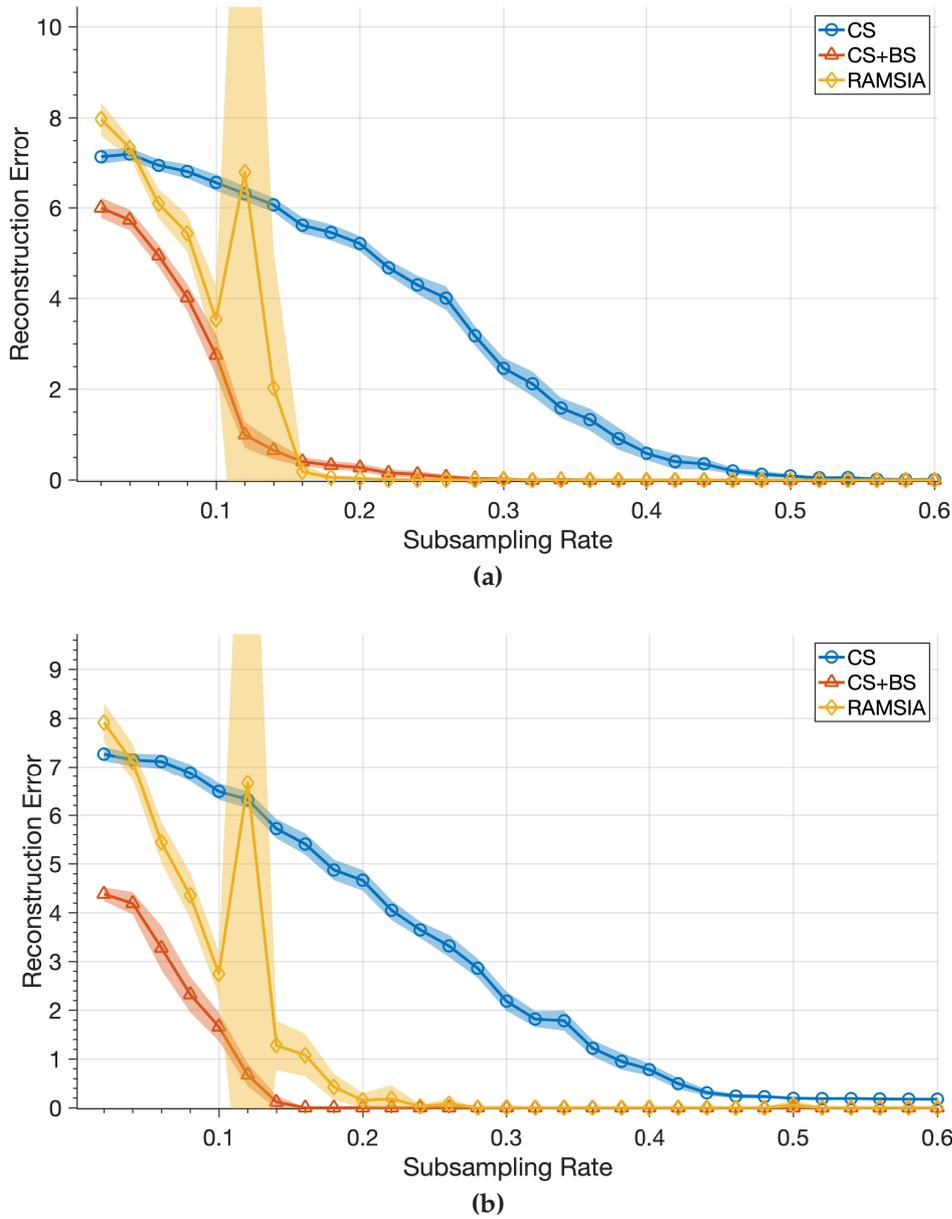


FIGURE 3.19: Mean reconstruction error and confidence intervals for increasing subsampling rates obtained with CS, CS with background subtraction and RAMSIA. (a) $k_x \leq k_z$ and (b) $k_x \geq k_z$.

Figures 3.19 (a) and (b), both depict the mean reconstruction error obtained after the reconstruction using a common CS basis pursuit algorithm (SPGL1) (the blue curve), using SI with the coherent background subtraction technique (the red curve) and using the RAMSIA algorithm (the yellow curve). In (a), we first create the ground truth vector x

with $k_x = 10$ nonzero elements. Subsequently, we create the SI z by replacing 10 random elements of x with random nonzero values. This means that the number of nonzero elements in x (k_x) is lower or equal to the number of non-zero elements in z (k_z). In (b) we create the inverse case where $k_x \geq k_z$. From this experiment, we can firstly easily conclude that both the coherent background subtraction approach and the RAMSIA algorithm are much more robust against subsampling compared to the common CS basis pursuit reconstruction. We further note that in (a), where k_x is smaller or equal than k_z , the RAMSIA algorithm reaches a mean reconstruction error close to zero at a lower subsampling rate compared to the coherent background subtraction approach. We observe the inverse behavior in (b), where k_x is greater or equal than k_z . The difference in reconstruction performance between these two reconstruction approaches in (a) and (b) can be attributed to the fact that both techniques minimize the ℓ_1 -norm of the difference between the two vectors, whereas only the algorithm with the weighted SI minimizes the ℓ_1 -norm of x and will thus be affected by its sparsity.

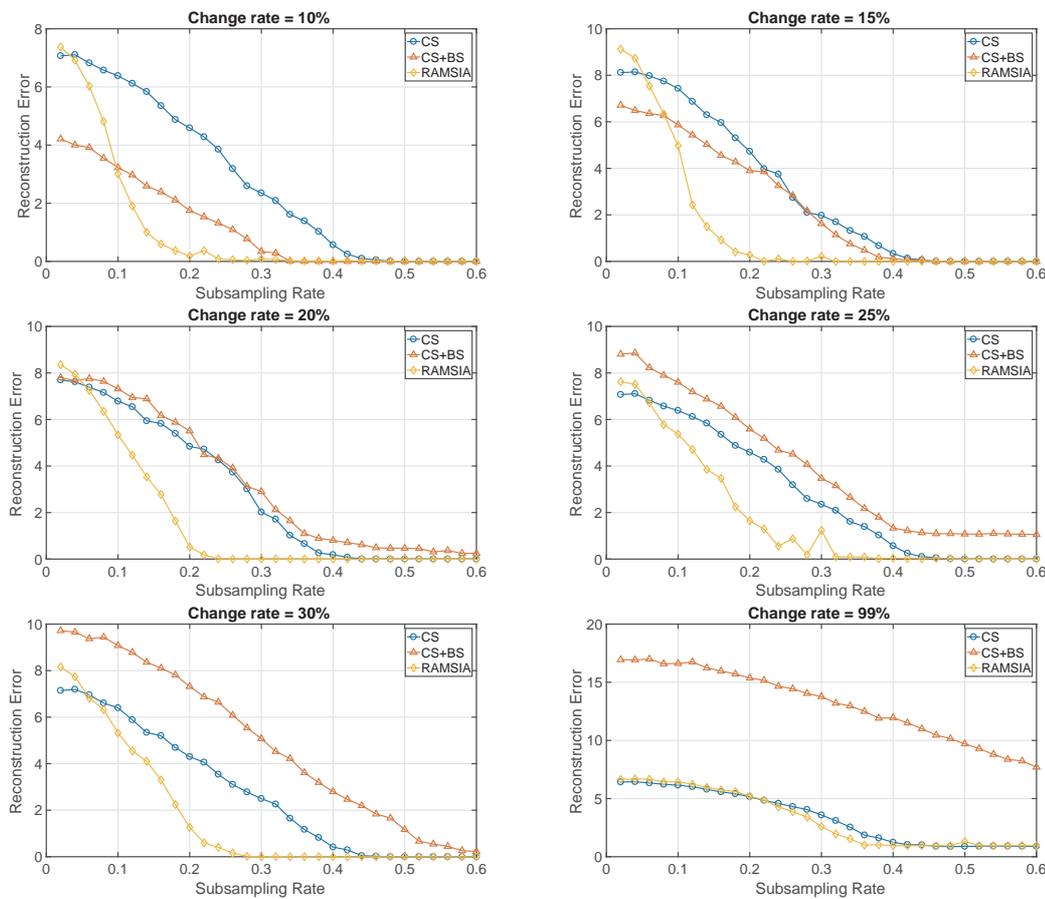


FIGURE 3.20: Mean reconstruction error for increasing subsampling rates obtained with CS (the blue curve), CS with coherent background subtraction (the red curve) and RAMSIA (the yellow curve) using a single SI with different quality levels (change rates: 10%, 15%, 20%, 25%, 30% and 99%).

We will now evaluate the robustness of the three CS reconstruction approaches against poor quality SI. These tests are relevant since we assume a high degree of similarity between the SI and the signal to reconstruct prior to the measurement. In the general case one is not able to accurately estimate this similarity since the signal to reconstruct is unknown. We start again by creating a random ground truth vector x with $k_x = 10$. We then create six SI vectors z with different change rates (ratio of random elements of x which are changed to a random nonzero value) equal to 10%, 15%, 20%, 25%, 30% and 99%. The vector x is again reconstructed using three CS reconstruction approaches: (1) a common CS basis pursuit algorithm (in blue), (2) CS combined with coherent background subtraction (the red curve) and (3) the RAMSIA (the yellow curve). Figure 3.20 depicts the mean reconstruction error over 32 executions of this experiment for increasing subsampling rates and using the six different SIs separately.

We firstly can conclude that the RAMSIA algorithm outperforms the two other approaches in each of the tests. The coherent background subtraction approach has a positive impact on the reconstruction performance up to a change rate equal to 20%. For higher change rates, we see that adding the SI using the coherent background subtraction approach deteriorates the reconstruction quality. The RAMSIA algorithm on the other hand, continues to improve the robustness against subsampling by filtering out the relevant information, hidden inside the SI and by neglecting the components of z possessing a low degree of similarity with x . Even for the extreme case where we change 99% of the elements of x to construct the SI z , RAMSIA does not have a negative impact on the reconstruction performance. The poor quality parts of the SI can not pollute the reconstruction thanks to the excellent performance of the intra-SI weights w_i .

We will now add multiple SIs with varying quality and evaluate the robustness of the coherent background subtraction technique and the RAMSIA algorithm. For these experiments, x has a length of 200 and contains 20 nonzero elements. Three different SIs are used for the reconstruction of x :

- SI1: A single SI with a change rate equal to 5% compared to x .
- SI2: Two SIs, with change rates equal to 5% for both SIs.
- SI3: Two SIs, with change rates equal to 5% and 20%.

For the coherent background subtraction approach, the background is obtained by calculating the mean of the different SIs. The results of these experiments are depicted in Figures 3.21 (a) and (b). For the coherent background subtraction approach in Figure 3.21 (a), we observe first of all, that adding SI1 and SI2 has the same positive impact on the reconstruction performance. The curves obtained for SI1 and SI2 are perfectly identical. Whereas, the enhanced performance is reduced for SI3 by adding SI with a lower degree of similarity. In (b) we observe that for SI1 (the red curve) the reconstruction performance is enhanced by adding SI1 compared to the common CS reconstruction without SI (the blue curve). When adding a second SI with the same change rate, the reconstruction performance is even further increased (which was not the case for the background subtraction approach). Finally, for SI3, we see that the reconstruction performance is better, even compared to adding SI1. The RAMSIA algorithm, firstly, favors the impact of the high quality SI with the inter-SI weights β_j and secondly filters the relevant information out from the poor quality SI thanks to the intra-SI weights w_{ji} . We can conclude from these experiments that the RAMSIA algorithm is immune for poor quality SI, which is

not the case for the coherent background subtraction approach. Adding SI by using the RAMSIA algorithm can only have a positive impact, and should thus be added for the reconstruction of x , regardless of the SI quality.

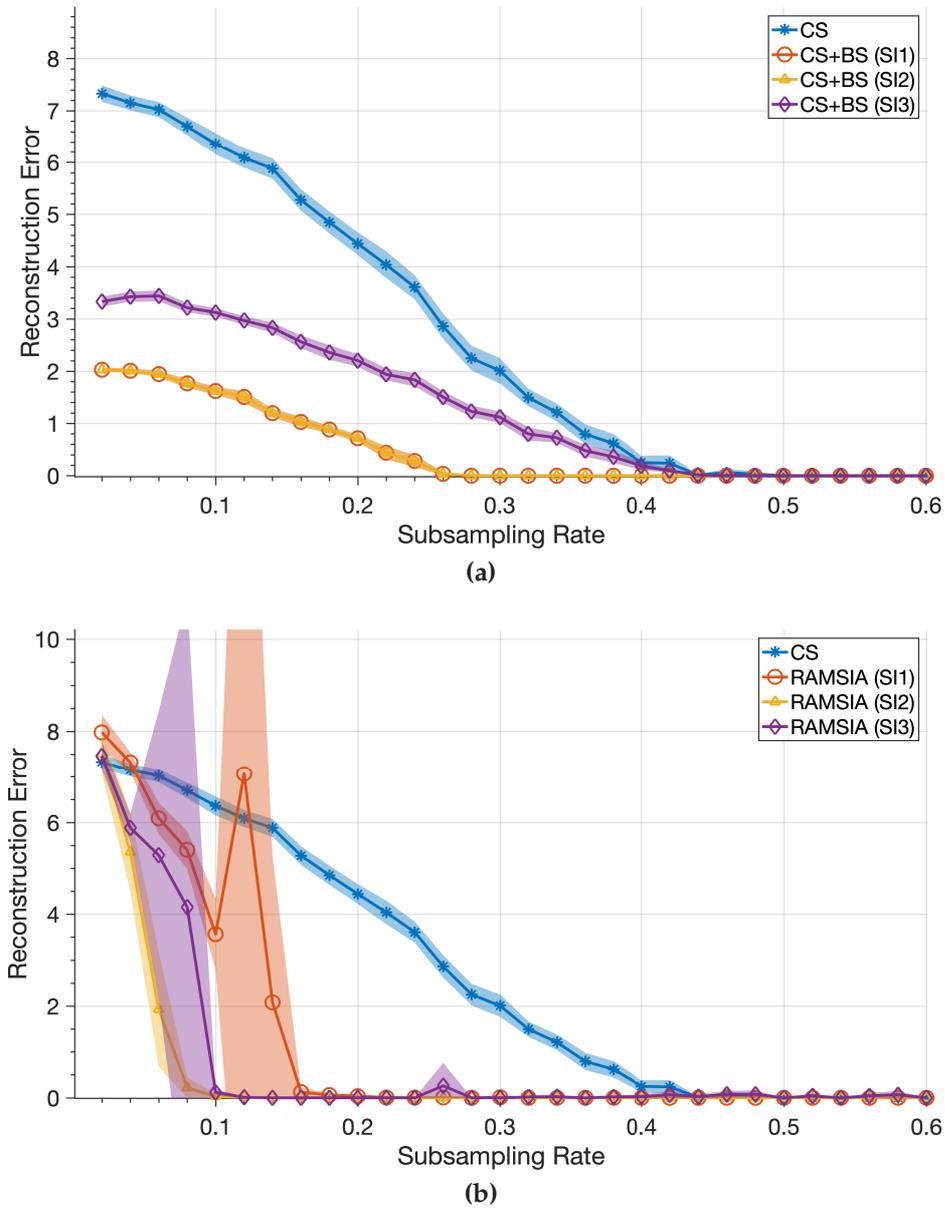


FIGURE 3.21: Mean reconstruction error and confidence intervals for increasing subsampling rates obtained with (a) CS (the blue curve) and CS with background subtraction and (b) CS (the blue curve) and RAMSIA for multiple SIs with varying change rates. (SI1) $J = 1$ (change rate 5%); (SI2) $J = 2$ (change rates = 5% and 5%); (SI3) $J = 2$ (change rates = 5% and 20%).

3.6.2 Heterogeneous Side Information

In the previous section we evaluated a technique based on a weighted n - ℓ_1 minimization approach for adding SI into the CS minimization. A high similarity between the SI and the actual signal to reconstruct is a key requirement to improve the robustness against subsampling. In other words, the SI must be obtained through a measurement performed with the same equipment under the exact same conditions of a very similar signal. In this section we will introduce and evaluate a strategy for adding heterogeneous SI. That is SI obtained from a measurement of a similar signal but with a different sensor or under different conditions. The explained methodology is based on the work presented in [Deligiannis et al.2017]. In that paper, the authors develop a source separation technique to separate an X-ray scan of the double-sided painted Ghent Altarpiece, by using photographs as heterogeneous SI. The strategy can be subdivided into two parts: First, a coupled dictionary learning method is developed. Second, the image separation is performed using the coupled dictionaries learned during the first part.

Coupled dictionary learning

Suppose the measurements y and x obtained with two different sensors:

$$\begin{aligned} y &= \Psi^c z \\ x &= \Phi^c z + \Phi v, \end{aligned} \quad (3.48)$$

where z is the signal with common components measured by the two sensors and v is the signal with the innovation components. The components of v contain the underlying information that is only measured by one of the two modalities. Ψ^c and Φ^c are the coupled dictionaries describing the linear measurement using the two measurement modalities of the information z sensed by both sensors. The matrix Φ describes the linear measurement of the innovation vector v . The three dictionaries can be learned through a training process, by solving following minimization problem:

$$\begin{aligned} \min_{\Phi^c, Z, \Psi^c, V} \quad & \frac{1}{2} \|Y - \Psi^c Z\|_F^2 + \frac{1}{2} \|X - \Phi^c Z - \Phi V\|_F^2, \\ \text{s.t.} \quad & \|z_t\|_0 \leq s_z, \quad \|v_t\|_0 \leq s_v, \quad \forall t = 1, \dots, T. \end{aligned} \quad (3.49)$$

with $\|M\|_F$ the Frobenius norm of the matrix $M \in \mathbb{C}^{m \times n}$, defined as:

$$\|M\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}. \quad (3.50)$$

The T columns of the matrices $Y \in \mathbb{C}^{N \times T}$ and $X \in \mathbb{C}^{N \times T}$ in (3.49) are populated with T corresponding training measurements of length N using both modalities. The minimization problem (3.49) searches a solution for the three dictionaries Ψ^c , Φ^c and Φ using the training measurements, while assuring a minimum sparsity smaller than or equal to s_z and s_v for the vectors populating respectively the matrices $Z \in \mathbb{C}^{N \times T}$ and $V \in \mathbb{C}^{N \times T}$.

The work in [Deligiannis et al.2017] proposes an algorithm for solving the minimization problem in (3.49). Each iteration k in the algorithm is subdivided in two steps: (1) A sparse coding step solving the following t parallel problems with fixed dictionaries:

$$(z_t^{k+1}, v_t^{k+1}) = \min_{z_t, v_t} \frac{1}{2} \left\| \begin{bmatrix} y_t \\ v_t \end{bmatrix} - \begin{bmatrix} \Psi^{c^k} & 0 \\ \Phi^{c^k} & \Phi^k \end{bmatrix} \begin{bmatrix} z_t \\ v_t \end{bmatrix} \right\|_F^2, \quad (3.51)$$

$$s.t. \quad \|z_t\|_0 \leq s_z \quad \text{and} \quad \|v_t\|_0 \leq s_v.$$

These minimizations problems are solved, using a Modified version of the Orthogonal Matching Pursuit (MOMP) algorithm. (2) A dictionary update step, solving following problems with fixed matrices Z and V :

$$\min_{\Psi^c} \frac{1}{2} \left\| Y - \Psi^c Z^{k+1} \right\|_F^2 \quad \text{and}, \quad (3.52)$$

$$\min_{\Phi^c, \Phi} \frac{1}{2} \left\| X - [\Phi^c \Phi] \begin{bmatrix} Z^{k+1} \\ V^{k+1} \end{bmatrix} \right\|_F^2.$$

Let us now evaluate this approach on synthetic data:

1. First, we choose two measurement modalities: (1) One modality described by a Fourier measurement matrix $\bar{\Psi}_c$ and, (2) one measurement modality defined by random Gaussian matrices $\bar{\Phi}_c$ and $\bar{\Phi}$.
2. We then create T vectors \bar{z} and \bar{v} with respectively s_z and s_v randomly chosen nonzero components.
3. The T signals \bar{z} and \bar{v} are then synthetically measured using the expressions in (3.48). The resulting measurements y_t and x_t , with $t = 1, \dots, T$, form the columns of the matrices X and Y and will serve as the training data for the next step.
4. The training data, obtained in the previous step, are used to solve the minimization problem in (3.49), giving us the learned coupled dictionaries Ψ^c , Φ^c and Φ and the sparse matrices Z and V . We note that the dictionaries Ψ^c , Φ^c and Φ and the sparse matrices Z and V , do not need to be the same as $\bar{\Psi}^c$, $\bar{\Phi}^c$, $\bar{\Phi}$, \bar{Z} and \bar{V} , to have successfully learned coupled dictionaries.
5. In a first approach for evaluating the performance of the coupled dictionary learning, we follow the methodology used in [Aharon et al.2006] and [Deligiannis et al.2017], by calculating the dictionary identifiability (further in this section the dictionaries will be evaluated after using them for performing the CS reconstruction with heterogeneous SI). The dictionary identifiability is calculated for each of the three dictionaries separately and is equal to the ratio of successful recovered columns. A column i of the learned dictionary Γ_i is considered to be successfully recovered if its distance compared with the closest column in the initial dictionary $\bar{\Gamma}_j$ is smaller than 0.01. The distances between the columns are calculated as ($|\cdot|$ denotes the absolute value):

$$1 - \left| \bar{\Gamma}_j^T \Gamma_i \right| \quad (3.53)$$

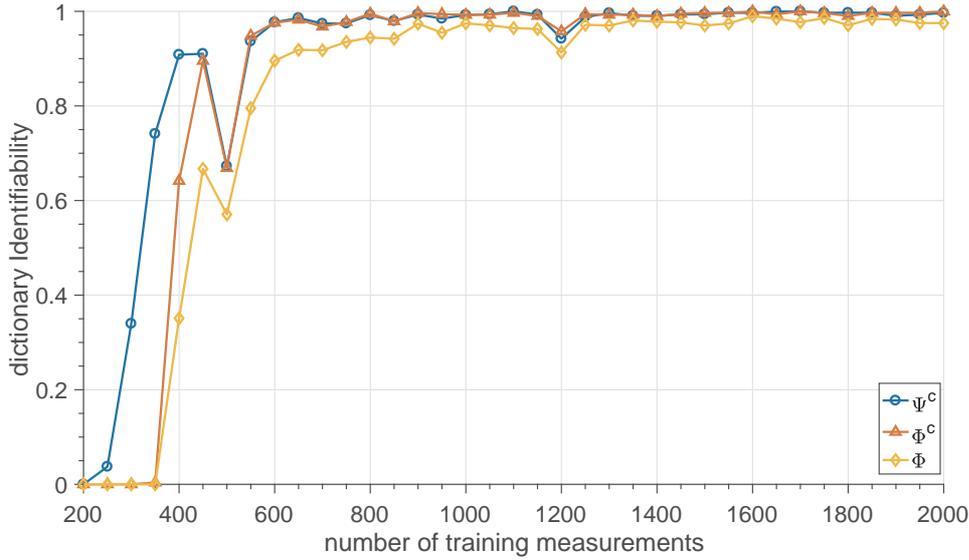


FIGURE 3.22: Dictionary identifiability of the coupled dictionaries Ψ^c , Φ^c and Φ learned from an increasing number of training measurements.

Figure 3.22 depicts the mean dictionary identifiability obtained over 64 independent coupled dictionary learning experiments for an increasing number of training measurements. The training measurements are obtained from simulated measurements of signals of length $N = 100$ and with sparsity rates $s_z = 5$ and $s_v = 3$. The initial dictionary $\bar{\Psi}^c$ is a Fourier matrix, $\bar{\Phi}^c$ and $\bar{\Phi}$ are chosen to be random Gaussian matrices. We observe that an identifiability very close to 1 is obtained with around 800 training measurements for Ψ^c and Φ^c . On the other hand, the dictionary Φ is learned more slowly and needs thus more training measurements.

CS with heterogeneous SI

Once the coupled dictionaries Ψ^c , Φ^c and Φ are learned, they can be used to reconstruct a signal with heterogeneous SI. Suppose that the coupled dictionaries are known and that a signal has been sensed using a first measurement modality, resulting in a measurement x . Suppose further, that a new measurement has to be performed using a different measurement modality. Thanks to the coupled dictionaries, the signal can be reconstructed from far fewer samples (measurements y) with a CS algorithm, by solving following minimization problem:

$$\min_{z,v} \left\| \begin{bmatrix} z \\ v \end{bmatrix} \right\|_1 \quad \text{s.t.} \quad \begin{bmatrix} y \\ x \end{bmatrix} = \begin{bmatrix} \Psi^c & 0 \\ \Phi^c & \Phi \end{bmatrix} \begin{bmatrix} z \\ v \end{bmatrix} \quad (3.54)$$

This approach is evaluated on synthetic data by using following methodology:

1. First we create the vectors \bar{z} and \bar{v} with respectively s_z and s_v nonzero elements. These two vectors will be the ground truth for the evaluation of this approach.

2. Next, we synthetically measure the signals \bar{z} and \bar{v} by applying the initial measurement matrices $\bar{\Psi}^c$, $\bar{\Phi}^c$ and $\bar{\Phi}$, resulting in measurements y and x .
3. Then, we randomly subsample the measurement y at a subsampling rate n/N .
4. Using the learned coupled dictionaries Ψ^c , Φ^c and Φ , we construct the vectors z and v , from the subsampled measurement y and the hetero-SI measurement x .
5. Finally, we reconstruct the full measurement $y = \hat{y}$ and the original vector $\bar{z} = \hat{z}$. The reconstruction is evaluated by calculating the relative reconstruction error: $\frac{\|\bar{z} - \hat{z}\|_2}{\|\bar{z}\|_2}$.

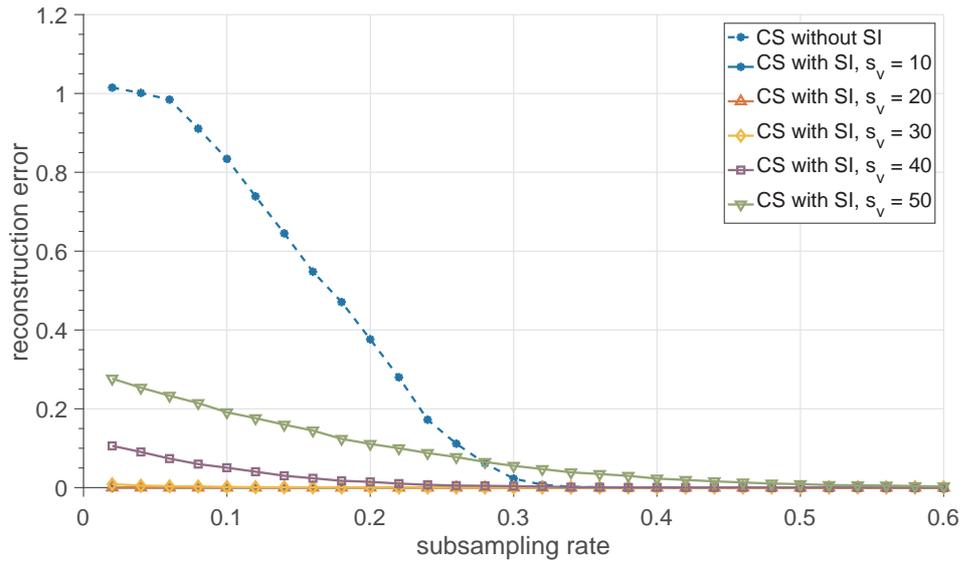


FIGURE 3.23: Reconstruction error for increasing subsampling rates when reconstructing the original signal with CS without SI (dashed line) and with hetero-SI (full line) for different values of s_v .

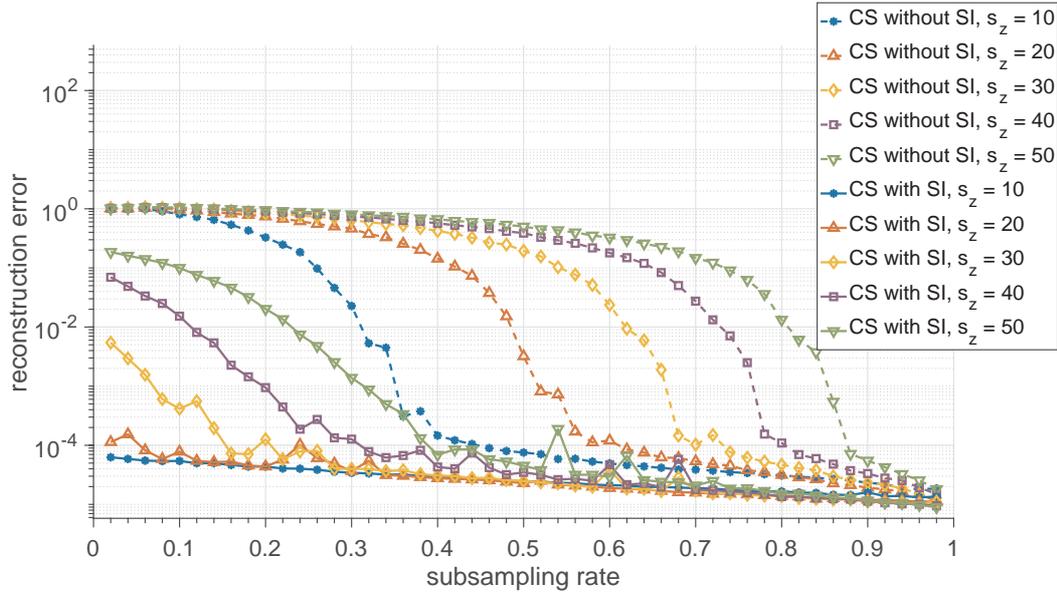


FIGURE 3.24: Reconstruction error for increasing subsampling rates when reconstructing the original signal with CS without SI (dashed line) and with hetero-SI (full line) for different values of s_z .

Figures 3.23 and 3.24 depict the results obtained from the experiments following the described methodology. The initial measurement matrix $\overline{\Psi}^c$ is again a Fourier matrix, $\overline{\Phi}^c$ and $\overline{\Phi}$ are chosen to be random Gaussian matrices. The dictionaries are learned from 2000 learning measurements. Figure 3.23 shows the reconstruction error when reconstructing a signal of length $N = 100$ and sparsity $s_z = 10$. The dashed plot is obtained when reconstructing the signal (with $s_z = 10$) using CS without SI, whereas the full lines correspond to the mean reconstruction errors over 100 executions when adding heterogeneous SI. The experiment is repeated for different values for s_v . We see that CS with heterogeneous SI clearly outperforms the CS reconstruction without SI as long as s_v stays below 40. For s_v equal to 50, we measure an increased reconstruction error between the subsampling rates equal to 0.28 and 0.50 for CS with heterogeneous SI when compared to CS without SI.

Figure 3.24 depicts the results obtained using the same dictionaries and same reconstruction methodology as for Figure 3.23, but this time s_v is chosen equal to 10 and we evaluate the mean reconstruction error for different values of s_z . Compared to the results obtained without SI, the approach using heterogeneous SI largely performs better for all values of s_z .

The mean reconstruction errors obtained after subsequently combining the coupled dictionary learning and the CS reconstruction with heterogeneous SI are depicted in Figure 3.25. First the coupled dictionaries are learned using a different number of learning measurement $T = 500, 750, 1000, 1250$ and 1500 . Then, a subsampled Fourier measurement of a random signal of length $N = 100$ (with 10 nonzero elements) is performed. From

this subsampled measurement, we reconstruct the original signal using the SI obtained from a random Gaussian measurement ($s_v = 10$) and the learned coupled dictionaries. Finally, we calculate the relative reconstruction error and we repeat this experiment 64 times in order to obtain a mean reconstruction error. From this experiment we can conclude: (1) The reconstruction is extremely robust and almost immune for subsampling. (2) The number of training measurements has an important impact on the reconstruction quality over all subsampling rates. This experiment gives a first idea on the order of magnitude for the number of training experiments needed.

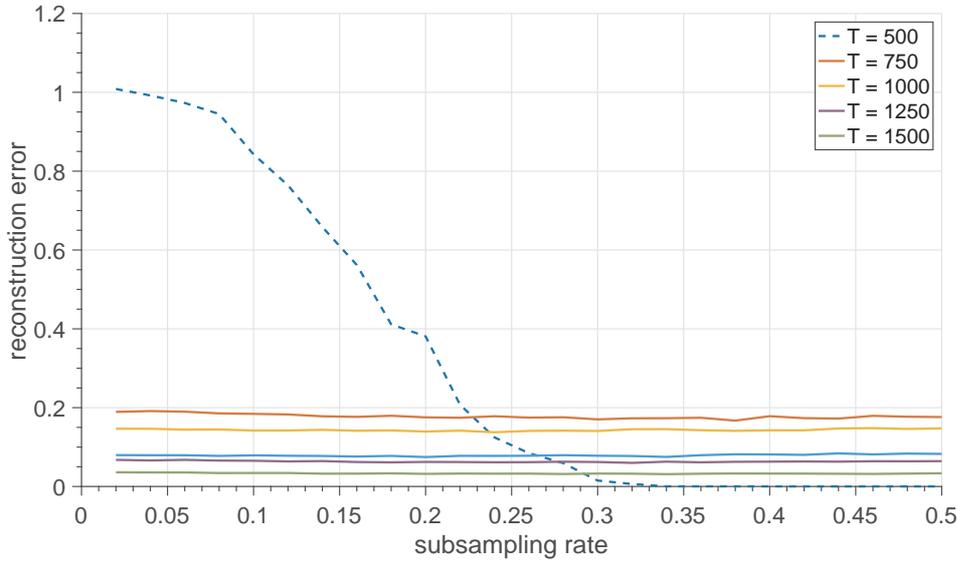


FIGURE 3.25: Reconstruction error for increasing subsampling rates when reconstructing the original signal with CS without SI (dashed line) and with hetero-SI (full line) using the learned dictionaries from different numbers of training measurements (T).

CS with heterogeneous SI for change detection

Until now, we assumed that the sensed signal did not change between the two heterogeneous measurements. In other words, the sensing modality described by Φ^c and Φ measures respectively z and v , whereas the modality described by Ψ^c measures exactly the same z . We will now suppose that the original signal changes between the two heterogeneous measurements by adding a vector δ , with s_δ nonzero elements, to z in the set of measurement equations:

$$\begin{aligned} y &= \Psi^c(z + \delta) \\ x &= \Phi^c z + \Phi v. \end{aligned} \quad (3.55)$$

This problem can occur with Non-Destructive Testing (NDT) measurements, where one tries to detect anomalies inside an object. Suppose for example that a critical part was tested using a first measurement methodology and later on, when a hidden defect appears, by a second methodology. The proposed approach will allow us to detect the

anomaly that is present in the second reconstructed signal using the heterogeneous side information from the first measurement. This can be achieved by solving following minimization problem:

$$\min_{z,v,\delta} \left\| \begin{bmatrix} z \\ v \\ \delta \end{bmatrix} \right\|_1 \quad \text{s.t.} \quad \begin{bmatrix} y \\ x \end{bmatrix} = \begin{bmatrix} \Psi^c & 0 & \Psi^c \\ \Phi^c & \Phi & 0 \end{bmatrix} \begin{bmatrix} z \\ v \\ \delta \end{bmatrix}. \quad (3.56)$$

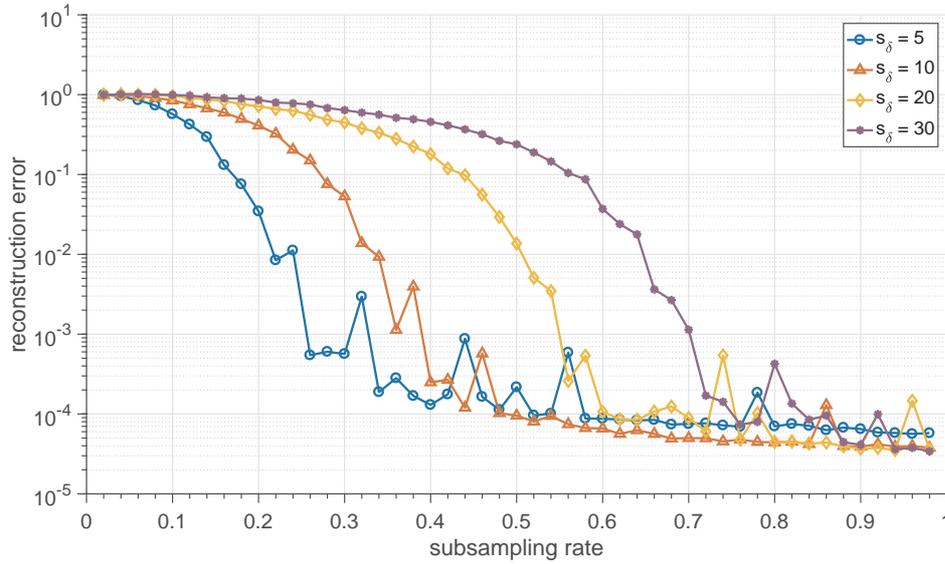


FIGURE 3.26: Reconstruction error for increasing subsampling rates when reconstructing the difference between two heterogeneous measurements using CS with heterogeneous SI.

Figure 3.26 shows the mean reconstruction error for reconstructing the signal δ , containing the change between the two heterogeneous measurements. For this experiment we first learn the coupled dictionaries from 2000 training measurements using a Fourier measurement matrix on the one hand and a random Gaussian matrix on the other hand. Using these dictionaries, we solve the minimization problem (3.56) from: (1) a measurement y , subsampled with increasing subsampled rates (horizontal axis), of a vector \bar{z} of length 100 with 10 nonzero elements to which we add the vector $\bar{\delta}$ with different numbers of nonzero elements (5, 10, 20 and 30) and (2) a heterogeneous measurement of \bar{z} and innovation vector \bar{v} ($s_v = 10$). The results depicted in Figure 3.26 proof firstly that this allows to perform change detection using heterogeneous SI and, second, the vector δ , describing the change, can be reconstructed from a subsampled measurement.

3.7 Conclusions

We started this chapter by evaluating if Compressed Sensing is a good candidate for the reconstruction of signals from subsampled SFCW or SAR measurements. The obtained

results, represented under the form of phase diagrams, showed that CS allows to take far fewer samples than predicted by the Shannon-Nyquist theorem. The minimum number of samples needed is unknown before the start of the measurement, since the lower bound depends on the unknown sparsity of the signal to reconstruct. In Section 4, we deduced and tested a method to estimate the needed number of samples without having access to the ground truth for SFCW measurements. We further showed that this estimation can be obtained from a small number of cross-validation measurements. In section 5, we explained how the signal processing tools SVD and RPCA can be applied to further improve the reconstruction performance when using CS. In section 6 we compared and evaluated different techniques for adding homogeneous SI to the CS minimization problem. The CS with multiple weighted SI was the best in all the tests. Finally, we explored a way to add heterogeneous SI and showed how to perform change detection using heterogeneous SI. In the next chapters, some of these techniques will be evaluated on real data.

Chapter 4

Non-Destructive Testing of polymer 3D-printed objects

4.1 Introduction

Additive manufacturing (AM) also known as 3D-printing, builds objects by joining materials one layer at a time, usually building them from the bottom up, leaving little to no waste [Hausman and Horne2014]. This revolutionary technique was developed back in the early 1980s and commercialized in the late 1980s [Hassen and Kirka2018]. Since then, the additive manufacturing market is exponentially growing. In 2016 the AM industry, consisting of all AM products and services, exceeded \$6.1 billion and is predicted to reach \$26.2 billion by 2022 [Wohlers et al.2016]. The interest for this manufacturing approach can be explained by its many advantages compared to traditional methods [Campbell et al.2011]: (1) Increased part complexity: the freedom to design parts is nearly unrestricted, allowing complex shapes and geometries that can be very difficult or even impossible to realize by the traditional methods such as molding, turning, milling and drilling. (2) Digital design and manufacturing: AM is a computer-controlled process which requires a low level of operator expertise and interactions. The physical part is created directly from a standardized digital file (.STL). (3) Instant production on a global scale: the digital file can be sent to any printer on the globe, or even in space, that can manufacture any product. (4) Waste reduction: In contrast to traditional subtractive manufacturing processes, such as machining, only the material needed for the part is used in production. There is virtually no waste. The 3D-printing technology can be subdivided into three categories based on the printed material:

1. Polymer materials. These can be preformed materials in powder, filament or sheet form. Alternatively, the material can be a photo sensitive resin [Stansbury and Idacavage2016].
2. Metals. Common materials used are stainless steels, aluminium, nickel, cobalt-chrome and titanium alloys [Yang et al.2017].
3. Other materials such as ceramics, concrete, etc [Trombetta et al.2017, Gosselin et al.2016].

Polymers are the most used material (85%) compared to metals (14%) and other materials (1%) [Wohlers et al.2016]. In this chapter we will only concentrate on 3D printed parts from polymer materials.

The additive manufacturing process of a part starts with a computer-aided design (CAD) model. This model is then sliced into layers which are printed by the AM system. The different printing approaches are:

1. Selective Laser Sintering (SLS): The feedstock for this type of AM system are thermoplastics such as polyamide (PA) and polycaprolactone (PCL) in powder form [Hassen and Kirka2018]. The first 2D slice is fabricated by fusing the powder using a laser beam after that the powder was spread on a platform. The next slices are formed by repeating this process while the platform is moving downwards in between the printing of each slice. This type of 3D printers can fabricate objects with complex geometries with high stiffness and strength properties. Due to the high cost, these machines are only used for the production of small volumes and high quality parts, as for aerospace prototyping applications [Wang et al.2017].
2. Stereolithography is performed by using a similar printing strategy. In this case, the feedstock material is a photopolymer, acrylate or epoxy-based resin which hardens when exposed to an ultraviolet laser beam. The feedstock materials for this type of AM machines are limited and expensive but they allow to create parts with high resolutions [Hassen and Kirka2018, Hull1986].
3. Fused Deposition Modelling (FDM) is the most widely, and cost-effective 3D methodology. It can be used with a wide range of polymers in a filament form, such as: polylactide (PLA), Acrylonitrile Butadiene Styrene (ABS), nylon 6 (PA6), nylon 66 (PA66), polycarbonate (PC) and polyetherimide (PEI). The filament is deposited on the existing layers by a heated nozzle which follows a predefined path. These systems have a limited speed and scale but are also low cost [Love et al.2015].
4. Extrusion Deposition or Big Area Additive Manufacturing (BAAM) is a more recent system. This type of AM printing machines use a similar technique as FDM but uses pellets as a feedstock material, which speeds up the printing process at the cost of a poor surface quality [Hassen et al.2016].

The quality of the manufactured part does not only depend on the quality of the feedstock material or the printing technique. The quality of the printed objects can vary between different AM machines and even between objects manufactured by the same system. Defects, and printing flaws are common in AM. The bonding between the different layers is created by the entanglement of the polymer chains. In order to make this bonding possible the previous deposited layer must locally remelt. Two important parameters affecting the quality during the printing process are thus the deposition temperature and the layer time. A maladaptation of one or both of these parameters to the part to be printed can cause different defects. This is illustrated in Figure 4.1, where three different defects are depicted. These defects were detected in objects printed at CISS department of the Royal Military Academy with a Stratasys SST 768 3D printer using ABS filament as feedstock material. In Figure 4.1 (a), we see a delamination. A delamination can be provoked by a long layer time. If the time interval between the deposition of two successive layers is too large, the previous deposited layer can cool down below the glass transition temperature, defined as the temperature at which the transition occurs from viscous liquid into an amorphous solid [Forrest et al.1996]. This results in a weak bond at the interface between the two layers and the two layers may delaminate [Hassen et al.2016]. In Figure 4.1 (b), we see a crack. Micro-cracks can appear across the part and originate from the residual stress caused by the large temperature gradients over the printed part.

Figure 4.1 (c) depicts an example of distortion or warpage. The heating and cooling in the FDM process causes residual stress which can provoke a deformation of the part. In this example, a rectangle massive block (width = 5 cm) was printed and one of the sides of the block came out curved, with a deformation of more than 0.35 cm.

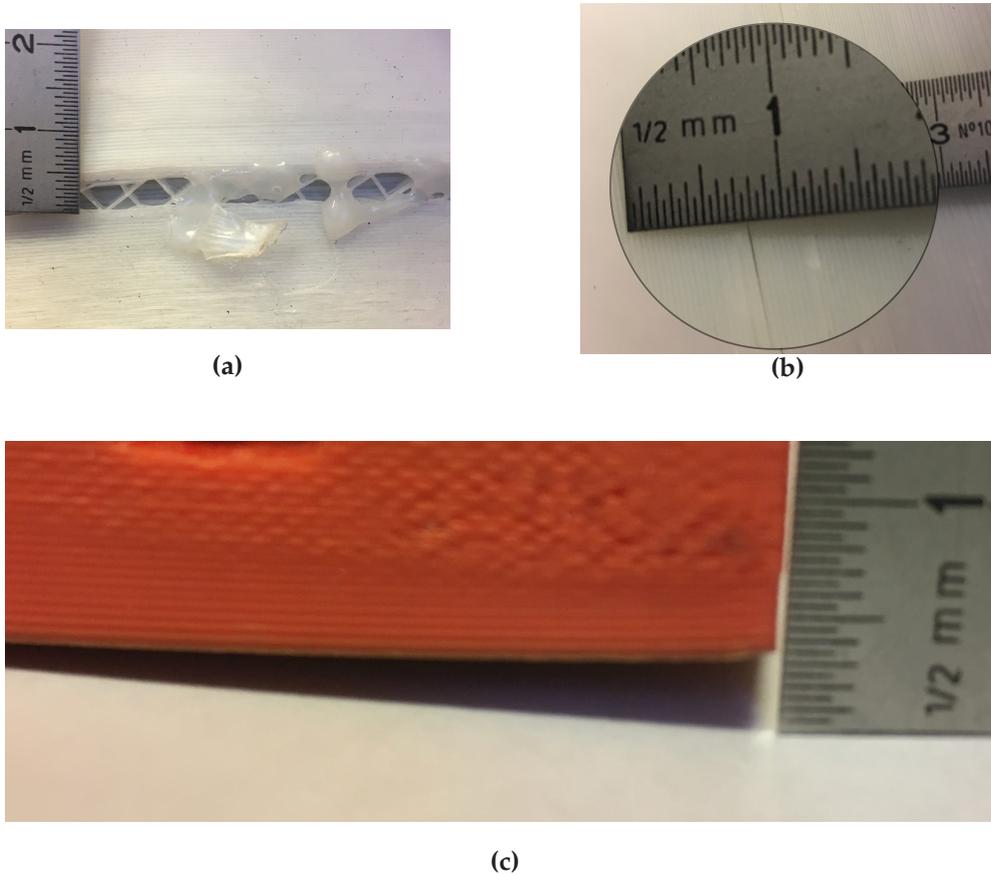


FIGURE 4.1: Examples of defects in 3D printed parts obtained with a single FDM machine.

The machine-to-machine as well as the day-to-day variability in quality of AM produced parts is one of the biggest challenges slowing down the transition from producing prototypes to the production of critical parts [Slotwinski2014, Everton et al.2016]. This is for certain the case in medical [du Plessis et al.2016], aeronautical [Seifi et al.2016], space [Waller et al.2015] and military applications. The NASA reports this lack of assurance of quality as one of the major obstacles for the acceptance of AM manufactured parts. More precisely, adequate non destructive evaluation processes are a universal need during process optimization, real-time process monitoring, finished part qualification and certification or in situ health monitoring [Waller et al.2015]. The today available NDT techniques for the quality control of additive manufactured parts are:

1. X-ray computed tomography is a method of forming three dimensional representations of an object by taking many x-ray images around an axis of rotation and using

algorithms to reconstruct a 3D model [Thompson et al.2016]. This NDT technique is the most frequent used technique for the Non-Destructive Evaluation (NDE) of AM parts due to its virtually unlimited in depth detection of anomalies. The main drawbacks of this NDT technique compared to other measurement techniques are [Thompson et al.2016, Lopez et al.2018]: (1) The measurements are time consuming and therefore not suitable for online evaluation during the production of the part. (2) The size of the area under test is limited. (3) This method is based on ionizing radiation which involves the need for health protective measures. (4) The systems are relatively costly.

2. Ultrasonic testing is a non-destructive evaluation technique which is broadly applied. Conventional pulse-echo systems can examine material properties including the microstructure and discontinuities in depth [Koester et al.2018, Rieder et al.2014]. However, this technique is not suited for high temperature inspection (typically $> 300^{\circ}\text{C}$). Moreover, this technology is not applicable to objects with rough surfaces or surfaces which are not locally planar, since it needs direct contact with the surface [Hassen et al.2016]. In general, extra machining, smoothing or polishing is needed in order to make ultrasonic testing possible. Laser ultrasonic testing offers a solution to these restrictions. These NDE systems direct a laser pulse on the surface of the part. The surface is locally heated which induces an ultrasonic pulse that propagates through the sample under test. A separate laser receiver detects the displacement of the surface caused by the backscattered pulse [Everton et al.2015].
3. Traditional IR thermography heats the part under test by a flash or quartz lamp. The induced heat flow is blocked by subsurface defects such as cracks or delaminations, causing a change in the transient heat flow response which can be observed with a IR sensor [Lu and Wong2018]. Infrared thermography is also well suited for the in-situ online monitoring during the layer by layer production of the part [Rodriguez et al.2015, Ding et al.2017].
4. Other NDT techniques include penetrant testing which is limited to surface testing and eddy current testing which is limited to conductivity materials [Lopez et al.2018].

Most 3D printing machines are open loop systems. An object is first completely manufactured and only after the production process a quality control is performed approving or rejecting the part. After production, measurement times are preferably kept as low as possible in order to reduce the down time due to the non destructive evaluation. Future NDT approaches for the quality control on 3D printed parts must preferably be integrated in a closed loop system. This implies that the NDT technique must be applicable during the manufacturing process in order to detect defects and make the adequate adaptations in time. In order to have a closed loop system, the system must be able to perform measurements at an equivalent rate compared to the printing speed (typical printing speeds are between 40 mm/s up to 150 mm/s). As well for open loop as for closed loop systems, the measurement time must be reduced to a minimum. Compressed Sensing can help to speed up the NDT measurement speed by lowering the number of required samples. Applying CS on ultrasound and X-ray tomographic imaging has extensively been covered in literature over the past years [Liebgott et al.2013, Schiffner and Schmitz2011, Yu and Wang2009, Xu et al.2012]. CS in combination with one of these imaging techniques is mostly applied in the field of medical imaging. In [Wagner et al.2012] for example, CS allows a nearly eight-fold reduction in sample-rate for cardiac ultrasound imaging.

In this chapter we explore a different NDT technique, based on mm-wave sensing [Ahmed et al.2012, Kharkovsky and Zoughi2007, Blitz2012, Kharkovsky et al.2006], applied on additive manufactured polymer objects. The proposed approach exploits the good penetration characteristics of microwaves through polymer materials. This characteristic enables the in-depth imaging of the inner structure of the sample. Moreover, in contrast to the NDT techniques based on X-rays, the radiation is non-ionizing and thus harmless. Compared to traditional ultrasonic testing, the presented approach operates from a stand off distance and can image objects with complex geometries or rough surfaces. The proposed microwave imaging approach needs to perform many measurements over a large bandwidth in order to produce high resolution images in the cross-track dimension. Each of these measurement must be executed at many sensor positions (in a SAR configuration) or object positions (in an ISAR configuration). Compressed Sensing techniques can be applied in order to obtain the same image quality from far fewer measurements. Recently, Helander et al. published a paper on mm-wave NDT of composite panels using CS [Helander et al.2017]. The work in this paper proposes to use CS for reducing the measurement time and improving the dynamic range. The proposed methodology starts with a source separation algorithm to separate the reflections originating from the defects from the rest of the backscattered signal. The separation is based on a Singular Value Decomposition (SVD). After the source separation, they use Compressed Sensing to reconstruct a mm-wave image of the defects. The technique was successfully tested on an industrially manufactured composite test panel.

The remainder of this chapter is structured as follows. In Section 2, we describe the mm-wave sensor, the setup and the measurements performed on a 3-D printed test object. The raw data obtained from these experiments will be exploited to evaluate different approaches for applying CS in the further sections of this chapter. In Section 3, we evaluate different approaches for applying CS on NDT ISAR data. We first compare the robustness against subsampling of a traditional ISAR reconstruction method with a Compressive Sensing reconstruction. We further evaluate the use of the 2-D DCT transform used as a sparsifying transform. If sparsity is not the only prior knowledge available, the complementary information, named Side Information (SI), can be added to enhance the reconstruction performance. We compare three different approaches for adding SI: (1) coherent background subtraction, (2) $\ell_1\ell_1$ -minimization and (3) weighted $\ell_1\ell_1$ -minimization. We end the section with the evaluation of CS with multiple weighted SI, using the real NDT data. The chapter ends with the conclusions obtained through the described experiments.

4.2 NDT experiment

In order to evaluate the applicability of CS for mm-wave NDT of polymer objects, a test set-up was built. The experiments were performed at the Department of Electronics and Informatics (ETRO) of the Vrije Universiteit Brussel (VUB). The sensor used for generating and acquiring the signals is based on a Millimeter Vector Network Analyzer (MVNA) from the company ABmm. The MVNA was connected to a single antenna: ABmm WR 15. The sensor generates Stepped Frequency Continuous Wave (SFCW) signals and measures the reflection parameter S_{11} :

$$S_{11}(f_p) := \frac{S_{Rx}(f_p)}{S_E(f_p)}, \quad (4.1)$$

where $S_{Rx}(f_p)$ is the received signal at frequency f_p and $S_E(f_p)$ is the emitted signal at the same frequency f_p . The total transmitted bandwidth is limited by the bandwidth of the antenna. The most important parameters of the sensor are enlisted in table 4.1.

Sensor parameter	Value
Starting frequency f_0	45 GHz
Bandwidth	30 GHz
Frequency step Δf	810.8 MHz
Scanning distance	0.30 m
Number of cross-range measurements	30
Number of frequencies	37
Aperture angle θ (-3 dB)	15°

TABLE 4.1

Figure 4.2 shows a schematic representation of the setup. The object under test is placed on a moving stage, 30 cm in front of the antenna. The moving stage is performing a horizontal displacement over 30 cm. During the horizontal displacement the moving stage stops 30 times. At each equidistant stop, a measurement is performed with the MVNA at 37 frequencies over a total bandwidth ranging from 45 GHz up to 75 GHz. Thanks to the horizontal displacement of the object under test, an ISAR image of a horizontal slice through the object can be constructed. The range resolution obtained with this sensor is defined by the total emitted bandwidth and is equal to:

$$\Delta r_{ra} = \frac{c}{2B}, \quad (4.2)$$

where c is the propagation speed of the electromagnetic waves through the medium with permeability μ and permittivity ϵ :

$$c = \frac{c_0}{\sqrt{\mu\epsilon}}. \quad (4.3)$$

The range resolution in vacuum will thus be equal to 5 mm. The objects under test are built from Acrylonitrile Butadiene Styrene (ABS) polymer which has a dielectric constant

around 3. Within the object under test the range resolution will thus be around 2.9 mm. Higher resolutions can be achieved by using larger bandwidths.

Figure 4.3 (a) shows a picture taken at the start of one of the experiments. The object under test is a block printed from ABS polymer material with dimensions: 18 cm \times 10 cm \times 5 cm (Figure 4.3 (b)). In total four separate measurements are performed consecutively:

- One measurement of the massive ABS block without synthetically created defects. This measurement will later in this chapter be used as Side Information (SI 1).
- Two measurements with one defect and with two defects respectively. The defects were created by drilling vertical holes with a radius of 5 mm through the massive block. SI 1 is the measurement of the block with one defect and SI 2 with two defects.
- One measurement of the ABS block with three holes (Figure 4.3 (b)). The raw data of this measurement will be undersampled and the mm-wave image of this object will be reconstructed using different CS reconstruction techniques.

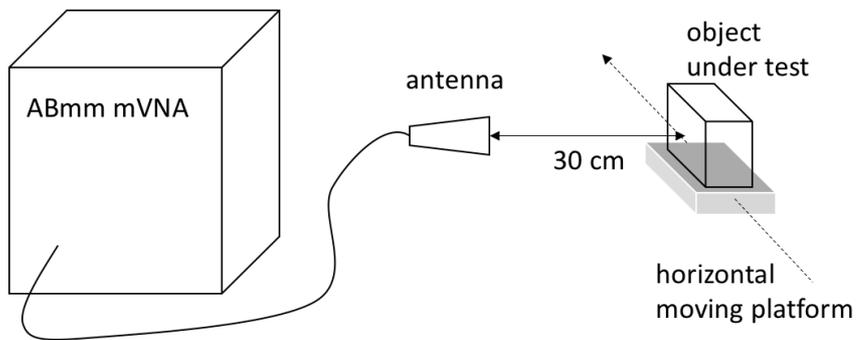
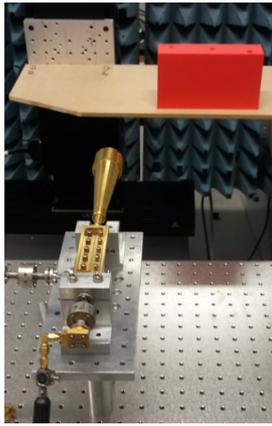
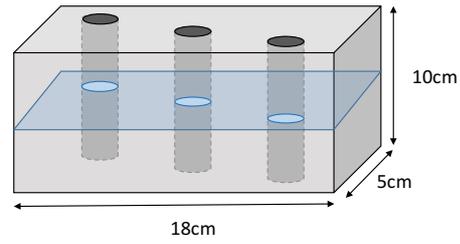


FIGURE 4.2: Scheme of the measurement setup for the ISAR NDT experiments.



(a)



(b)

FIGURE 4.3: (a) picture taken during a NDT measurement of the test object. (b) schematic representation of the test object with the three defects.

Figure 4.4 depicts the mm-wave images obtained from the ISAR measurements of the four objects under test. The four experiments are performed under similar conditions. The images are obtained from a CS reconstruction of the fully sampled ISAR data. These experiments allow clearly to image the front and backside of the object under test and to detect the different defects.

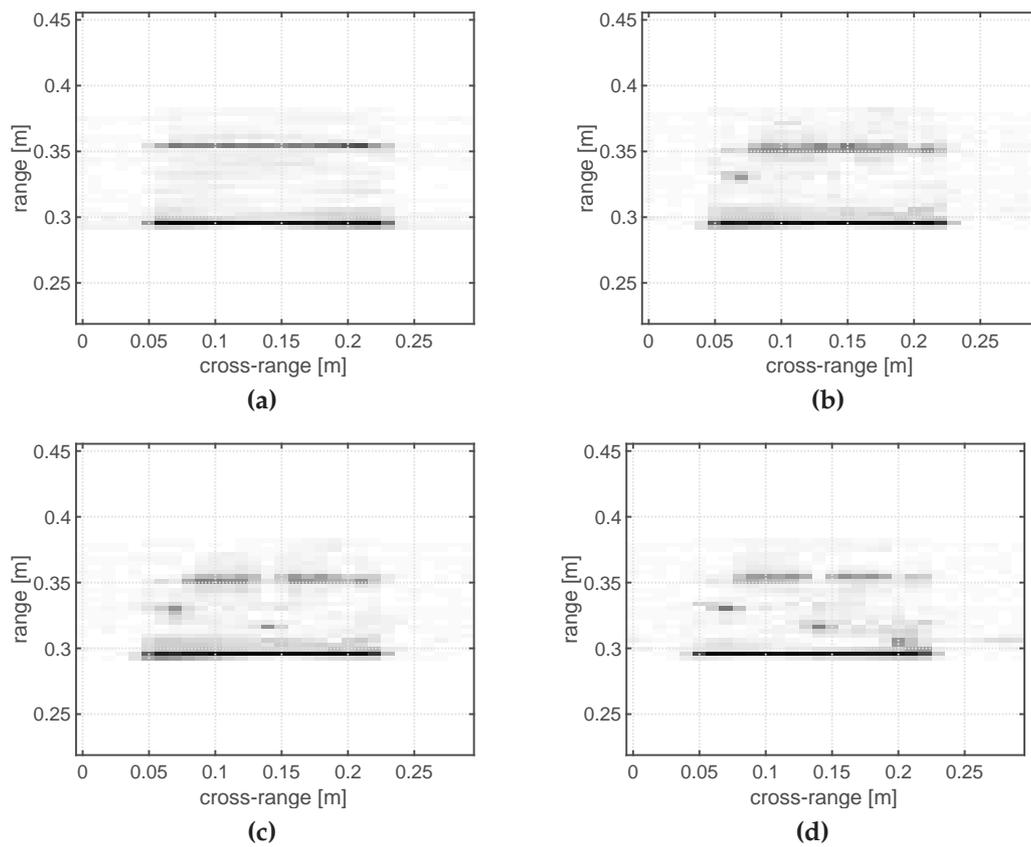


FIGURE 4.4: mm-wave ISAR images (horizontal cut) of the test object in (a) without defects, in (b) with 1 defect, in (c) with 2 defects and in (d) with 3 defects.

4.3 CS reconstruction

One of the key requirements for a NDT technique is the need for a small measurement time. Small measurement times will on the one hand reduce the off-duty time of a 3D printed part during a (recurrent) inspection or on the other hand, the measurement can be performed online during the production of the part. A reduction of the measurement time can be achieved by reducing the number of samples to be taken. In this section we evaluate the robustness against random undersampling of the raw data of the ISAR measurements obtained from the previous described experiments.

4.3.1 CS reconstruction versus traditional SAR reconstruction

Compressive Sensing is based on the assumption that the signal to reconstruct is sparse. The sparser the signal, the fewer samples needed to obtain a correct reconstruction of the signal. Figure 4.5 (a) depicts a histogram of the normalized pixel values of the mm-wave image of the ABS block with three defects. We observe that most of the pixels have a value close to zero. Instead of reconstructing the image directly, we can opt for reconstructing a sparsified transform of the image. Figure 4.5 (b) shows the histogram of the 2-D Discrete Cosine transform of the mm-wave image with the three defects. The DCT transformed signal has indeed more elements with a value close to zero: 310 pixels have a normalized value within the interval 0 - 0.01, whereas the original image has only 52 pixels with a normalized value in the same interval.

Figure 4.7 depicts the images obtained with:

1. A common range Doppler algorithm (left column) [Ozdemir2012]. These images are obtained in two steps: (1) A range compression: obtained by performing an IFFT of the raw data in the range dimension and (2) A cross-range compression: the range compressed data is therefore transformed to the range-Doppler domain and a range cell migration correction is performed before removing the cross-range chirp. Finally, the image is obtained after performing an IFFT in the cross-range dimension. We can clearly see that the range Doppler algorithm is not robust against subsampling. The obtained image is already deteriorated at a subsampling rate of 90% and at 70% the defects are drowned in the reconstruction noise.
2. A common CS reconstruction algorithm (middle column). The reconstruction is much more robust against subsampling. At a subsampling rate of 50%, it becomes difficult to detect the defects and at a subsampling rate of 40% the image is too much deteriorated to perform an evaluation of the object under test.
3. A common CS algorithm used to reconstruct the 2-D DCT transformed mm-wave image. The results are visually comparable to those obtained without sparsifying transform. The images look a bit blurred when compared to the images of the middle column.

Figure 4.8 shows a numerical evaluation of the reconstruction performance of the three approaches for increasing subsampling rates together with the symmetric 95% confidence intervals. For each subsampling rate, starting at a subsampling rate of 0.02 up to 1 the raw data are randomly subsampled before reconstructing the mm-wave image, using the

three different approaches. This test was repeated 100 times with each time a new independent random sampling of the raw data. The full curves in Figure 4.8 show the mean reconstruction error which was calculated as the distance between the reconstructed signal from the subsampled raw data and the signal obtained from the fully sampled data. The dashed lines indicate the standard deviation from the mean reconstruction error. We deduce following conclusions from this experiment:

- The reconstruction quality starts to deteriorate immediately when reducing the number of samples when applying a traditional (I)SAR reconstruction algorithm.
- As expected from the visual evaluation, the CS reconstruction is much more robust against subsampling.
- When using the 2-D DCT transform, we observe first a better reconstruction performance below a subsampling rate of 0.52 compared to the CS reconstruction without sparsifying transform. Between subsampling rates equal to 0.52 and 0.9, using the DCT does have a negative impact on the reconstruction performance.

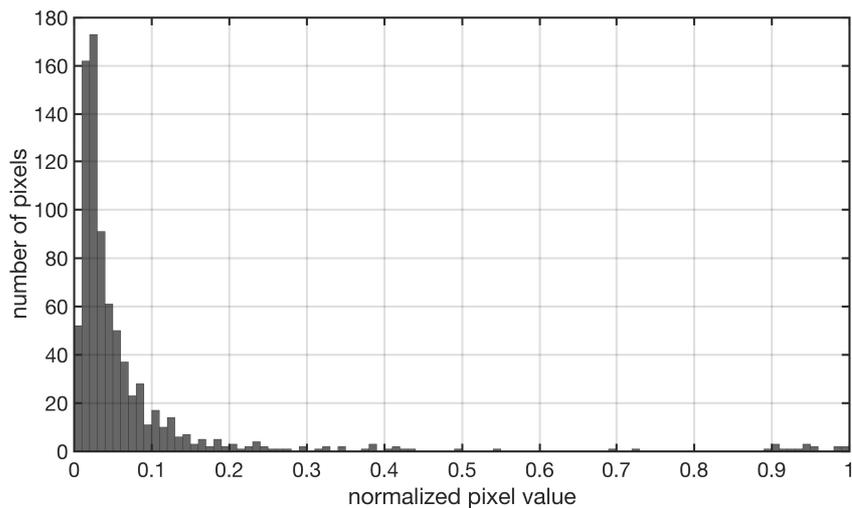
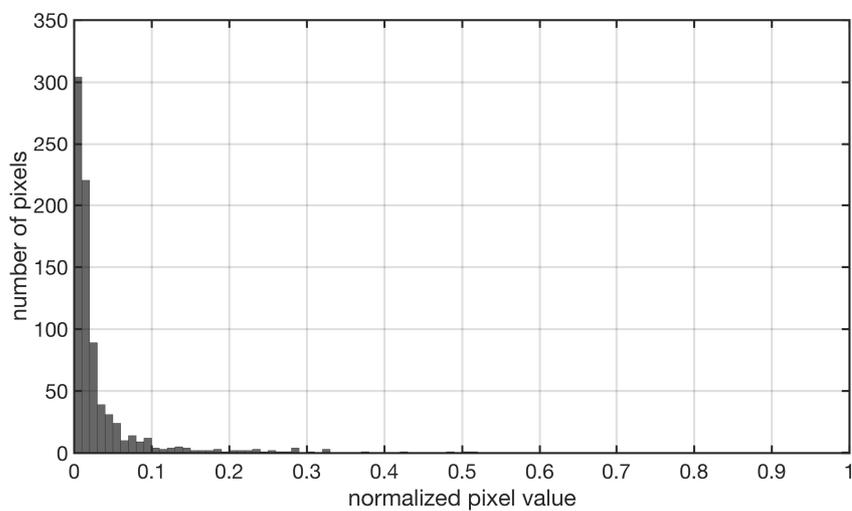
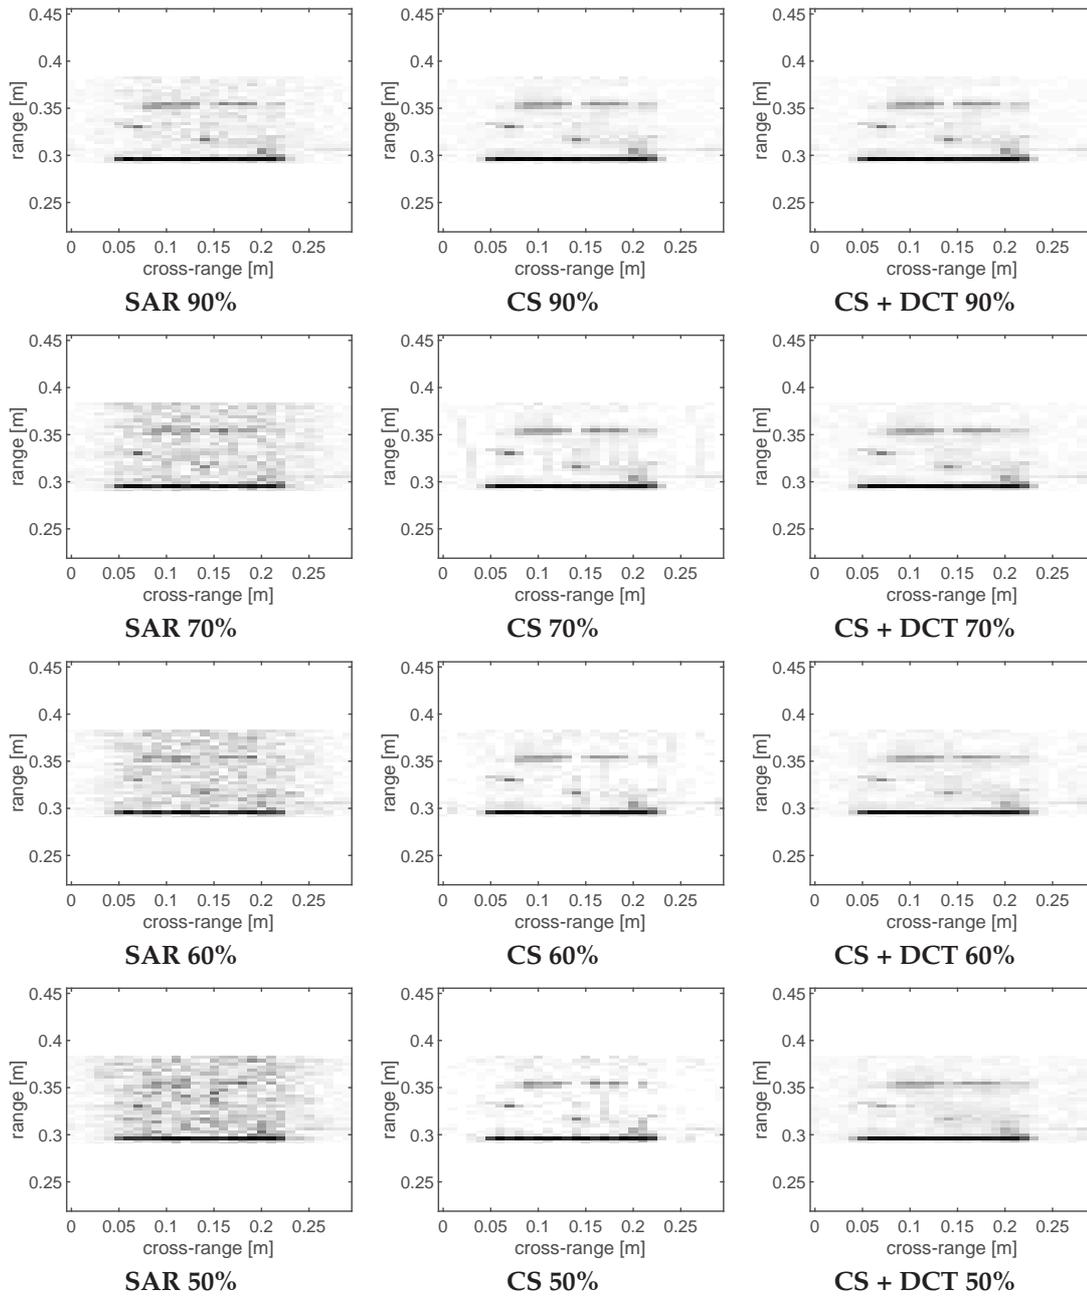
**(a)****(b)**

FIGURE 4.5: Histograms of the normalized pixel values. (a) Histogram of the ISAR image to reconstruct. (b) Histogram of the DCT of the ISAR image to reconstruct.



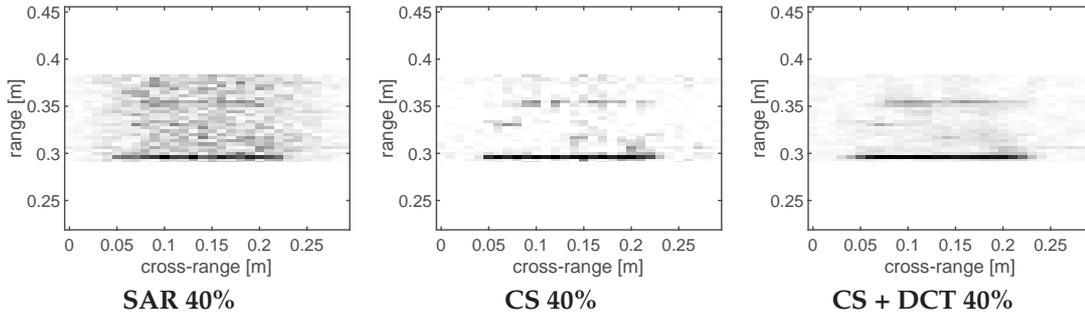


FIGURE 4.7: reconstruction for different subsampling rates. Left column: images obtained with traditional SAR algorithm, middle column: results obtained with CS and right column: results obtained with CS and the DCT transform.

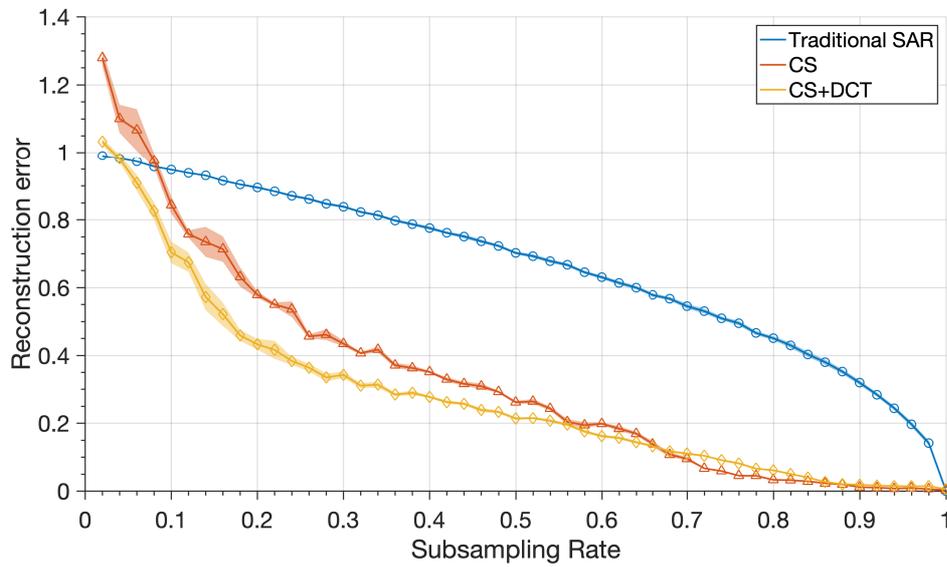


FIGURE 4.8: Mean reconstruction error and confidence intervals for increasing subsampling rates obtained with a standard SAR algorithm, with a CS algorithm (SPGL1) and with a CS algorithm with a DCT transform. The dashed curves are the corresponding standard deviations.

4.3.2 CS with side information

We will now reconstruct the NDT image with three defects exploiting side information. The side information that will be used for these experiments are the measurements of the samples without defects, with one defect and with two defects. We reconstruct the sample with three defects using three different strategies: (1) CS with Coherent background Subtraction (BS), (2) $\ell_1\ell_1$ -minimization and (3) CS with weighted Side Information (SI).

CS with Coherent background subtraction

Instead of reconstructing the SAR image of the sample with three defects, we reconstruct the difference between the SAR image of the sample under test and the SAR image obtained from a prior measurement of a similar object. The basic idea is that the difference will be sparser than the actual image. This hypothesis is true if the prior SAR image shares a high degree of similarity with the SAR image to reconstruct.

The experiment is performed as follows: First, the fully sampled data of the measurement of the sample with three defects is randomly subsampled. Second, the corresponding elements from the raw data of the SI is subtracted from these samples. Subsequently the SAR image of the difference between the two SAR images is reconstructed using a CS algorithm. Finally the SAR image of the samples with three defects is obtained by adding the SAR image of the SI to the reconstructed difference image. We performed this experiment for an increasing number of samples, starting at a subsampling rate (n/N) equal to 0.02 up to 1. The experiment was repeated 100 times for each subsampling rate with each time a new random sampling of the raw data. The results in Figure 4.9 show the mean distance over the 100 experiments between the reconstructed SAR image from the subsampled raw data and the SAR image obtained from the fully sampled data. This experiment was repeated with three different SI's:

- SI 1: the sample without defects (Figure 4.4 (a)): The ℓ_2 -norm of the difference between SI 1 and the SAR image to reconstruct is equal to 0.0805. We observe a clear improvement of the reconstruction quality over all the subsampling rates. A mean reconstruction error below 0.2 is achieved with a subsampling rate of 60% with CS without the use of SI, whereas this result is obtained with only 38% of the samples using the background subtraction approach with SI1.
- SI 2: the sample with one defect (Figure 4.4 (b)): The ℓ_2 -norm of the difference between SI 2 and the SAR image to reconstruct is equal to 0.1071. Surprisingly, SI 2 shares a decreased similarity with the SAR image to reconstruct compared to SI 1. This reduced SI quality results in a decreased reconstruction performance. Compared to the mean reconstruction error obtained without CS we can identify two zones: (1) below a subsampling rate of 40%, the BS with SI2 performs better than the reconstruction without SI whereas (2) for subsampling rates between 40% and 100% the SI deteriorates the reconstruction performance using the BS approach.
- SI 3: the sample with two defects (Figure 4.4 (c)): The ℓ_2 -norm of the difference between SI 3 and the SAR image to reconstruct is equal to 0.2408. The similarity for SI 3 is thus even further decreased compared to SI 1 and SI 2. The reconstruction error with SI 2 is larger than the reconstruction error without SI. Adding SI of this quality deteriorates the reconstruction quality over all the subsampling rates.

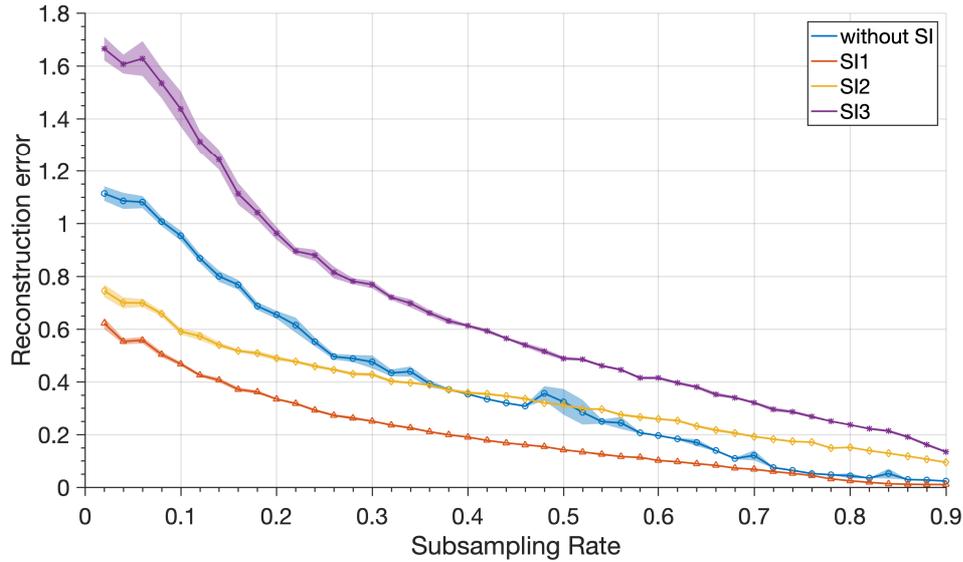


FIGURE 4.9: Mean reconstruction error for increasing subsampling rates obtained from the CS SAR image reconstruction with coherent background subtraction without SI, with SI 1, SI 2 and SI 3.

Coherent background subtraction can have a positive impact on the reconstruction quality by enhancing the robustness against subsampling. We also conclude from this experiment that the similarity between the SI and the SAR image to reconstruct must be high in order to obtain a better reconstruction performance. The coherent background subtraction reveals to be very sensitive to poor quality SI.

$\ell_1\ell_1$ - minimization

With this approach, the sum of the ℓ_1 -norm of the SAR image with the ℓ_1 -norm of the difference between the SAR image to construct and the SAR image of the SI, is minimized. In contrast to the coherent background subtraction approach, the sparsity of the difference between the SAR image to reconstruct and the SI is promoted during the reconstruction instead of reconstructing the difference. Figure 4.10 depicts the mean reconstruction errors for increasing subsampling rates. The experiment is again performed with the same three SI's:

- SI 1: The $\ell_1\ell_1$ -minimization using SI 1 as side information has a clear positive impact on the reconstruction quality over all subsampling rates. We measure a slight improvement compared to the coherent background subtraction approach: a mean reconstruction error equal to 0.2 is obtained with only 34% of the samples instead of 38%.
- SI 2: The same trend can be observed when using the SI 2. Again, the reconstruction performance is improved over all subsampling rates compared to the CS reconstruction without SI. Due to the decreased similarity, the positive impact of adding the SI

in the minimization is reduced. Compared to the coherent background subtraction approach, we note a clear reconstruction improvement.

- SI 3: We observe a deterioration of the reconstruction performance for subsampling rates between 34% and 100%. The negative impact is less important compared to the results obtained with the coherent background subtraction technique.

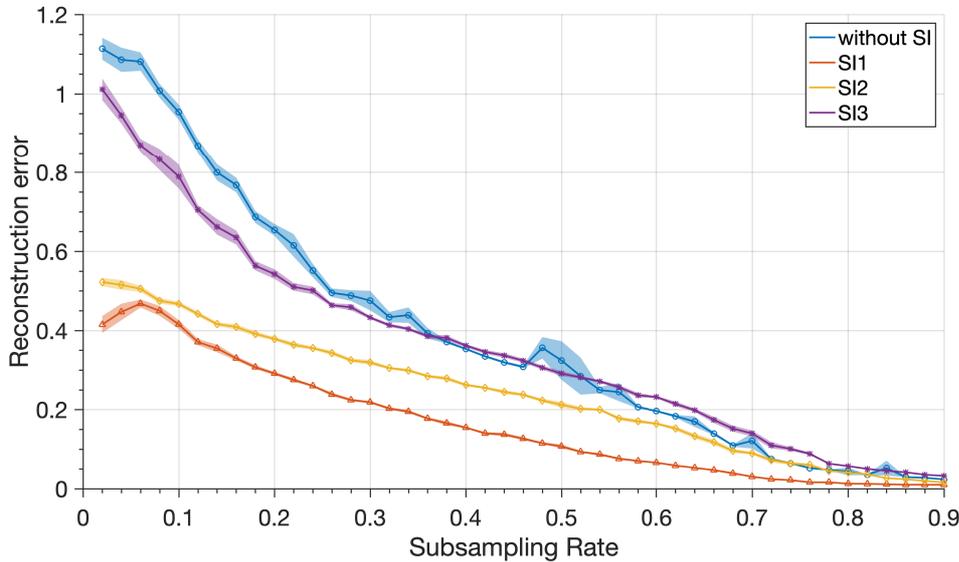


FIGURE 4.10: Mean reconstruction error for increasing subsampling rates obtained from the CS SAR image reconstruction with $\ell_1\ell_1$ -minimization without SI, with SI 1, SI 2 and SI 3.

From this experiment we can conclude that the $\ell_1\ell_1$ -minimization approach is less sensitive for poor quality SI compared to the coherent background reconstruction. This reduced sensitivity does not impact the positive improvement of adding high quality SI to the CS reconstruction. If the similarity between the SI and the SAR image to reconstruct decreases too much, the poor quality SI can have a negative impact on the reconstruction compared to the CS reconstruction without SI.

CS with Weighted side information

Adding SI as a background image for coherent background subtraction or by using the $\ell_1\ell_1$ -minimization approach can improve the reconstruction quality if the SI shares a high degree of similarity with the image to reconstruct. On the other hand, if the difference between the SAR image to reconstruct and the SI is too important, adding SI with one of these two approaches can have a negative impact on the reconstruction performance. The RAMSIA algorithm, explained in paragraph 3.6.1, offers a solution to prevent a deterioration of the reconstruction performance by poor quality SI. The side information is

added during the reconstruction of the SAR image from the undersampled raw data and the reconstruction is obtained by solving the following minimization problem:

$$\min_x \left\{ \frac{1}{2} \|Ax - y\|_2^2 + \lambda \sum_{j=0}^J \beta_j \|W_j(x - z_j)\|_1 \right\}, \quad (4.4)$$

where x is the mm-wave image to be reconstructed, y the undersampled SAR measurement, z_j are the SI's and z_0 is equal to 0. The impact of the SI during the minimization is regulated by two weighting parameters: β_j and W_j .

β_j are inter-SI weights. The RAMSIA algorithm allocates an inter SI-weight to the ℓ_1 -norm of x : β_0 and to the ℓ_1 -norm of the difference between x and the J SI's: $\beta_{1\dots J}$. The higher the weight attributed to the SI, the higher the impact of the SI in the minimization of the equation (4.4). The sum of all inter-SI weights is equal to 1. Figure 4.11 depicts the evolution of the inter-SI weights obtained with RAMSIA for increasing subsampling rates. The figure shows the weight for the reconstruction with a single SI. The test was performed with each of the three SI's separately. At very low subsampling rates, the weights vary severely. The weights start to stabilize from a subsampling around 30%. The inter-SI weight attributed to the SI 1 is relatively high and evolves finally towards a value equal to 0.6456. This means that the sparsity of the difference between x and the SI 1 has a larger impact on the minimization of equation (4.4) than the sparsity of the SAR image to reconstruct itself ($\beta_0 = 0.3544$). The inter-SI weight for SI 2 evolves towards a value equal to 0.5212 and towards 0.3508 for SI 3. The order of the values of the inter-SI weights is in concordance with the quality of the SI estimated by calculating the distance between x and the SI's (table 4.2).

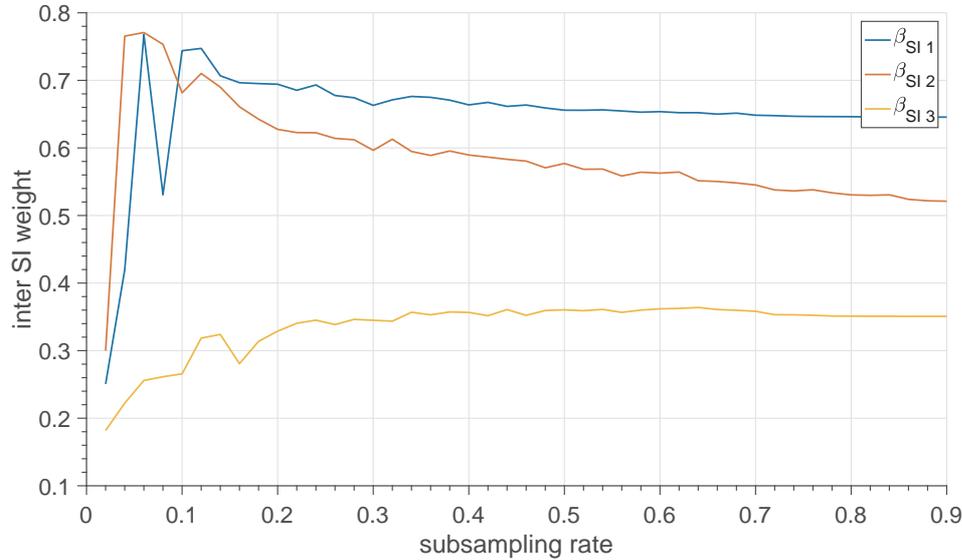


FIGURE 4.11: inter SI weights β_1 using a single SI (SI 1, SI 2 or SI 3), for increasing subsampling rates.

	SI 1	SI 2	SI 3
$\ x - SI\ _2$	0.0805	0.1071	0.2408
β_0	0.3544	0.4788	0.6492
β_{SI}	0.6456	0.5212	0.3508

TABLE 4.2

The quality of the SI does not only vary between the multiple SI's but can also change inside a single SI. This variation of similarity inside (intra) the SI's is countered with the intra-SI weights. $W_j \geq 0$ is a diagonal weighting matrix attributed to the SI z_j and whose diagonal elements ($w_{j1}, w_{j2}, \dots, w_{jn}$) are populated with the weights allocated to the different elements of z_j . The weights are iteratively obtained by the RAMSIA algorithm. Elements differing much from the corresponding element of the SI will receive lower weights. On the other hand, elements with values close to the value of the corresponding element of the SI will get a higher weight.

The influence on the reconstruction performance of the intra-SI weights is evaluated in Figures 4.12 and 4.13. Figure 4.12 (a) and 4.13 (a) depicts the difference in absolute value between the 810 elements (horizontal axis) of the SAR image to reconstruct and the elements of the SI's SI 1 and SI 3 respectively. We note that the images share a high degree of similarity since the graph is close to zero except for some peaks. When comparing SI 1 and SI 3, we observe that the peaks are much higher for SI 3. Figures 4.12 (b) and 4.13 (b) show the intra-SI weights attributed to the 810 elements of SI 1 and SI 3 respectively without subsampling. We observe that the graphs follow closely the inverse trends of the graphs in (a). Figures 4.12 (c) and 4.13 (c), finally show the intra-SI weights obtained with RAMSIA for a subsampling rate equal to 30%. Even at low subsampling rates the weights continue to follow the inverse trend of the difference in (a). This method of attributing weights to the different elements of the SI is thus robust against subsampling.

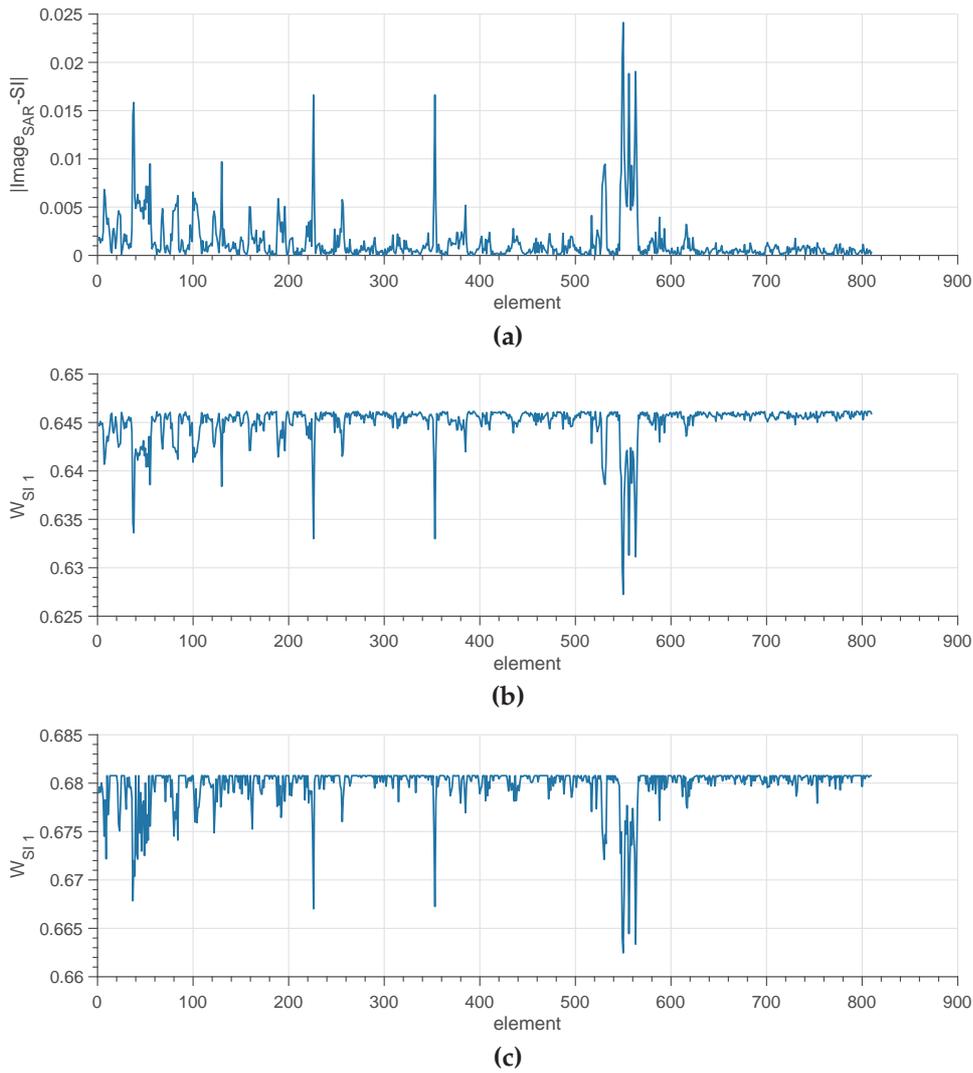


FIGURE 4.12: (a) Absolute value of the difference between the SAR image to reconstruct and the SI 1 over all the elements of the SAR image. (b) The intra SI weights obtained with RAMSIA for a subsampling rate equal to 90%. (c) The intra SI weights (SI = SI 1) obtained with RAMSIA for a subsampling rate equal to 30%.

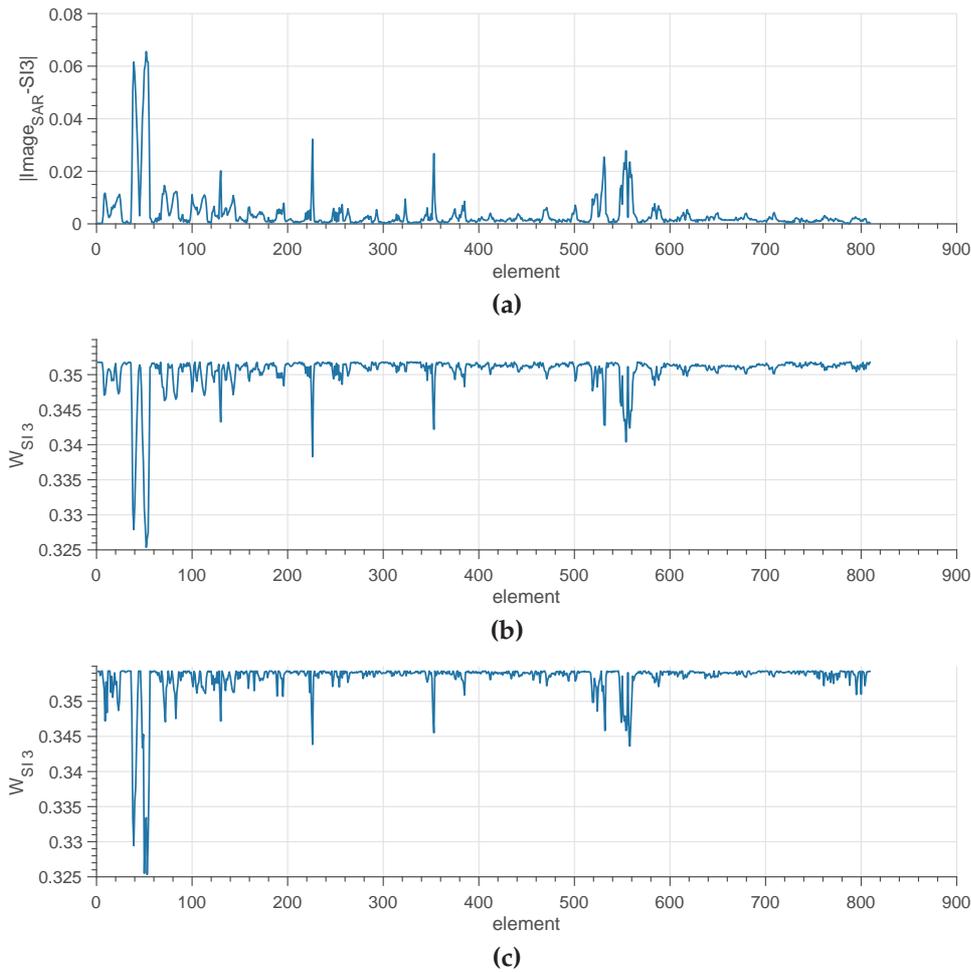


FIGURE 4.13: (a) Absolute value of the difference between the SAR image to reconstruct and the SI 3 over all the elements of the SAR image. (b) The intra SI weights obtained with RAMSIA for a subsampling rate equal to 90%. (c) The intra SI weights (SI = SI 3) obtained with RAMSIA for a subsampling rate equal to 30%.

Figure 4.14 depicts the mean reconstruction error for a subsampling rate equal to 0.02 and increasing with a stepsize of 0.02 (horizontal axis) obtained with the RAMSIA algorithm with a single SI. The experiment was performed 100 times for each of the subsampling rates and with each of the three SI's separately. Using SI 1 and SI 2, we obtain a very similar result as with the $\ell_1\ell_1$ -approach. For example, a mean reconstruction error equal to 0.2 is obtained at a subsampling rate of 0.32 with the RAMSIA approach using SI 1 whereas we obtained this result at a subsampling rate of 0.34 with the $\ell_1\ell_1$ -minimization technique using the same SI. Comparing the results obtained with RAMSIA and the other two techniques for SI 3, we observe that the mean reconstruction error does not pass the

reconstruction error for the CS reconstruction without SI. This was not the case for the coherent background subtraction or the $\ell_1\ell_1$ -minimization approach. Thanks to the weighting mechanisms, the reconstruction performance can not deteriorate when adding SI into the RAMSIA reconstruction. In other words: it is not necessary to estimate the quality of the SI in order to decide if the SI can be added into the CS reconstruction. Adding poor quality SI will be neglected thanks to the inter-SI weights. Moreover, the RAMSIA algorithm will take the similar parts of the SI in consideration for the reconstruction, while neglecting the elements of the SI which differ too much, thanks to the intra-SI weights.

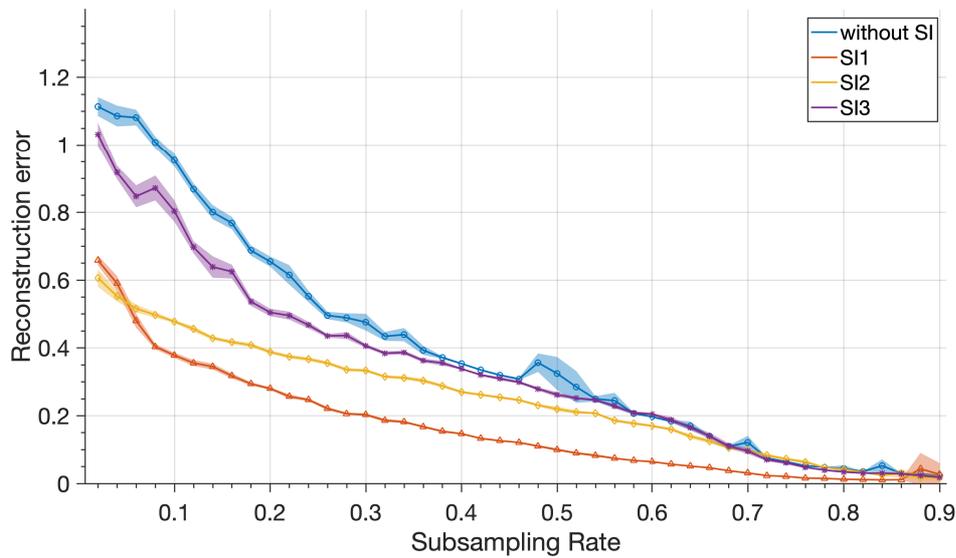


FIGURE 4.14: Mean reconstruction error for increasing subsampling rates obtained with the CS SAR image reconstruction with weighted SI: without SI, with SI 1, SI 2 and SI 3.

4.3.3 CS with multiple homogeneous weighted SI

In this section, multiple SI's will be exploited for the reconstruction of the subsampled SAR image. We compare three different approaches. A visual comparison of the reconstructed SAR images of the test object with three defects, for subsampling rates equal to: 60%, 40%, 30% and 20%, is depicted in Figure 4.15:

1. Coherent background subtraction (figures in the left column of Figure 4.15). In this case, each of the elements of the background is equal to the mean value obtained from the corresponding elements of the different SI's. If we compare the obtained images to those obtained with CS without SI (the middle column in Figure 4.7), we do not see a big improvement. At a sampling rate of 60% the defects are still clearly visible, whereas at a subsampling rate of 40% this is not the case anymore.
2. $\ell_1\ell_1$ minimization (figures in the middle column of Figure 4.15). A single SI is obtained by taking the mean value of each element over the different SI's. The reconstruction quality for lower subsampling rates is slightly better than the quality obtained with coherent background subtraction or without using SI.
3. CS with multiple weighted SI (figures in the right column of Figure 4.15). The SAR images obtained with weighted SI are clearly of a higher quality than the ones obtained with coherent background subtraction or $\ell_1\ell_1$ -minimization. In contrast to the previous approaches, the CS with multiple weighted SI technique allows to clearly detect the defects at a subsampling rate of 40% and even at 30% the defects are still visible. The detection of the defects starts to become difficult at a subsampling rate around 20%.

The same conclusions can be deduced from the graphs in Figure 4.16, where the mean reconstruction error, over 100 executions, is depicted for increasing subsampling rates. We observe that the mean reconstruction error using the background subtraction approach is even worse than the reconstruction without SI for subsampling rates above 52%. At lower subsampling rates, below 40%, the coherent background subtraction approach performs better than the CS reconstruction without SI. But, the defects are no longer visible at those low subsampling rates. The $\ell_1\ell_1$ -reconstruction technique performs better than the coherent background subtraction or the CS reconstruction without SI up to a subsampling rate equal to 72%. At subsampling rates between 70% and 90% both $\ell_1\ell_1$ -minimization and the CS reconstruction without SI perform evenly. Further, we clearly see that the RAMSIA algorithm outperforms all the other tested approaches over all subsampling rates.

The good reconstruction performance of the RAMSIA algorithm can be explained by the excellent impact of the weights. Figure 4.17 depicts the variation of the inter-SI weights β , where β_0 corresponds to the weight attributed to the ℓ_1 -norm of the SAR image to reconstruct and β_1 , β_2 and β_3 are the weights allocated to SI 1, SI 2 and SI 3 respectively. At subsampling rates below 20% the weights vary, but for subsampling rates above 30% the weights vary only smoothly. We further observe that the weight β_3 stays very low over for all subsampling rates. Indeed SI 3 is the poorest SI and thus will almost be neglected by the RAMSIA algorithm. Whereas SI 1, the SI which shares the highest similarity with the image to reconstruct, has the largest impact in the minimization equation (4.4).

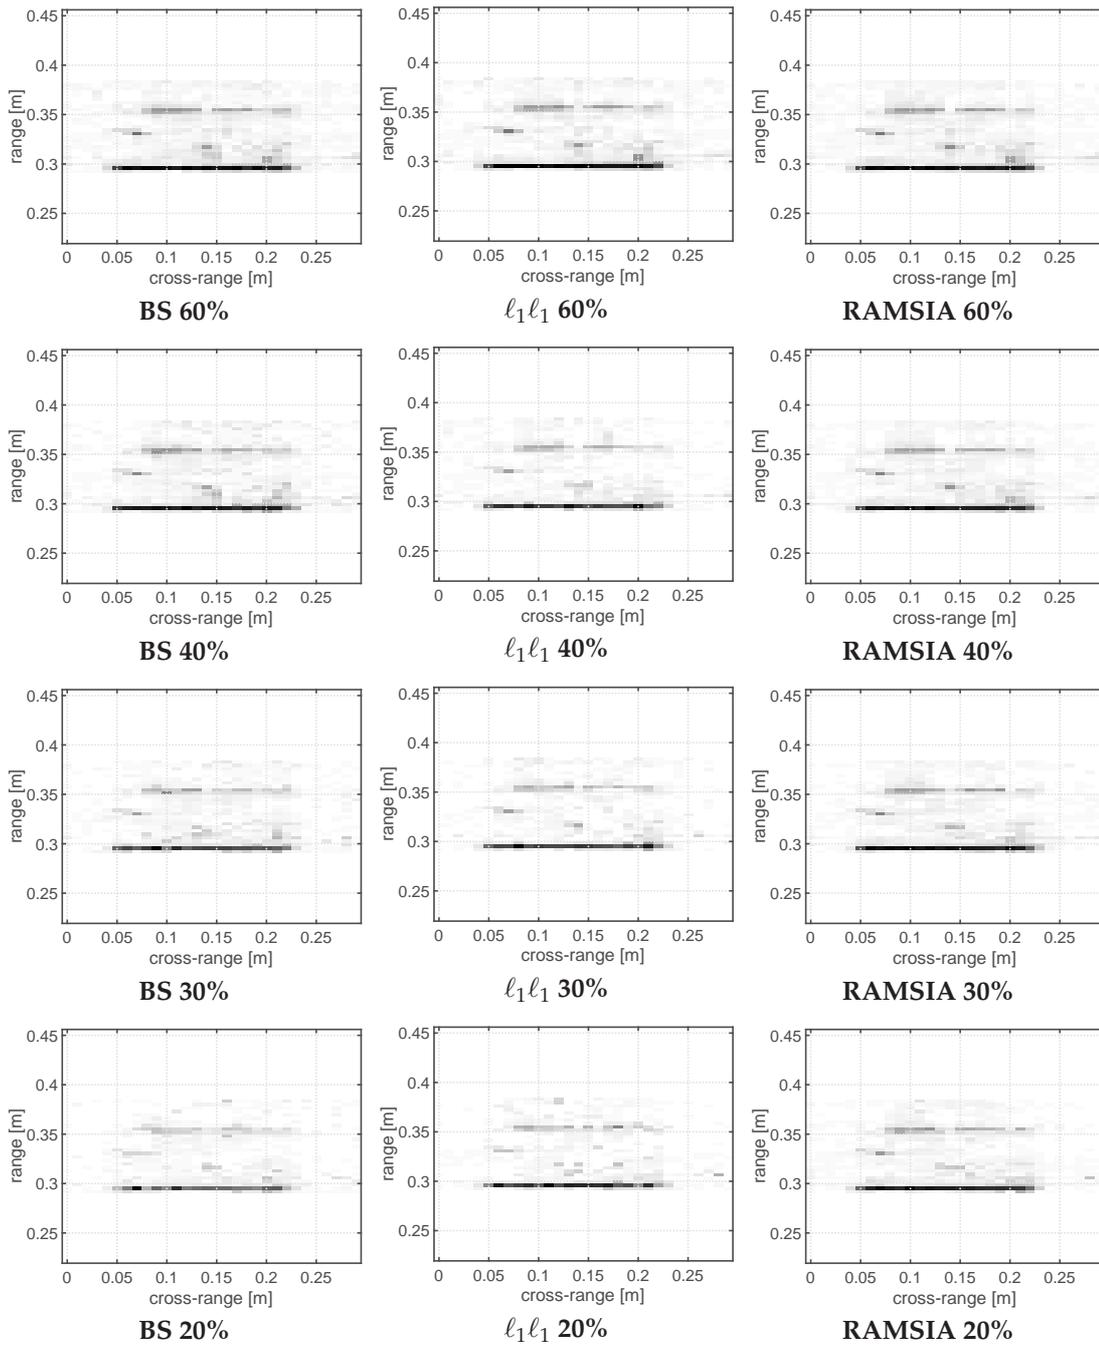


FIGURE 4.15: Reconstructed SAR images obtained with coherent background subtraction (left column), $\ell_1\ell_1$ -minimization (middle column) and RAMSIA (right column) for subsampling rates equal to 60%, 40%, 30% and 20%.

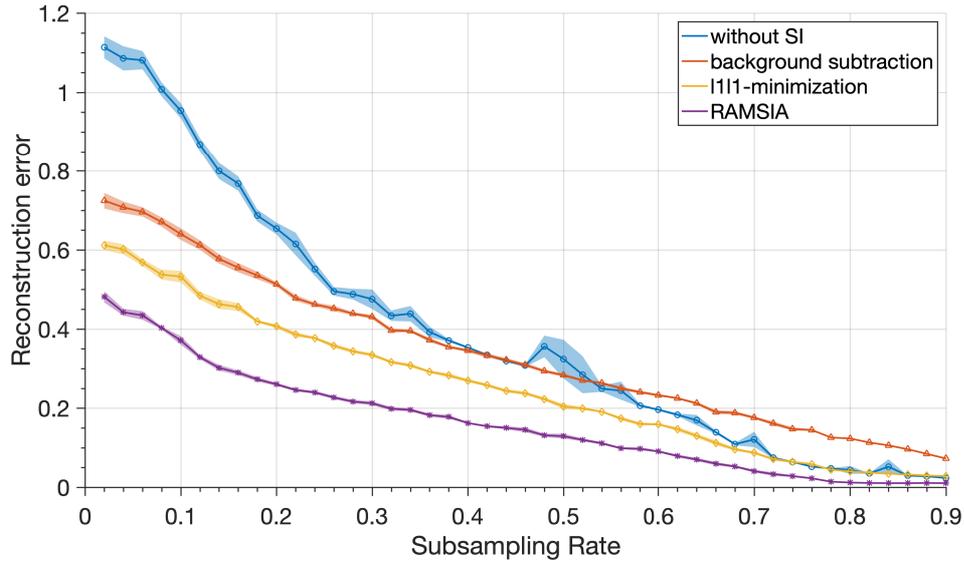


FIGURE 4.16: Mean reconstruction error for increasing subsampling rates obtained with coherent background subtraction, $\ell_1\ell_1$ -minimization and RAMSIA using the three SIs .

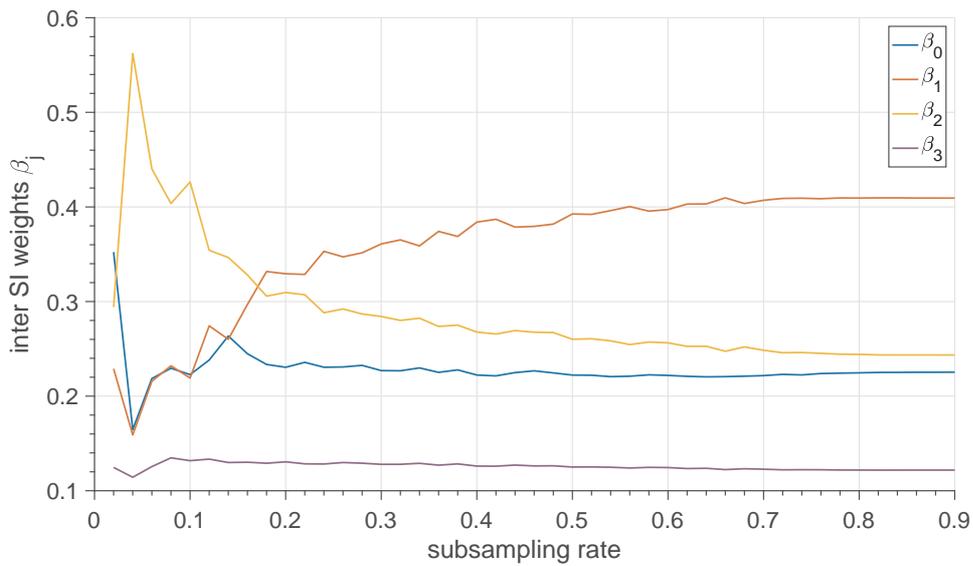


FIGURE 4.17: inter SI weights β_j for increasing subsampling rates obtained with the RAMSIA algorithm.

4.4 Conclusions

This chapter started with a small introduction on additive manufacturing and explained the need for adequate NDT techniques. In this chapter we proposed a mm-wave imaging technique and tested this technique on a 3-D printed object with artificially created defects. One of the drawbacks of this approach is the long measurement time and the high data volumes created during the measurement. This is a major issue, since the measurement speed is of high importance in NDT applications.

We demonstrated on the real data that Compressed Sensing allows to create high quality mm-wave images from severely subsampled raw data. We further tested the use of a 2-D DCT sparsifying transform. We observed that the augmented sparsity by the 2-D DCT did not have a positive impact on the reconstruction performance.

Often, sparsity is not the only prior knowledge that can be used to reduce the number of measurements. For the production of parts by additive manufacturing, a blueprint of the object to print does exist. From this blueprint a prediction can be made of the measured signal during the non-destructive evaluation of the part. Or, if multiple parts are produced, the mm-wave images should be identical except for printing errors or defects. Finally, in recurrent inspections, a measurement taken at a previous moment in time does exist. This supplementary information, known as Side Information can be exploited during the mm-wave image construction from the subsampled raw data. In this chapter we evaluated three approaches to add the SI: (1) coherent background subtraction, (2) $\ell_1\ell_1$ -minimization and (3) CS with weighted SI. The three techniques were tested using real data SI of different quality. From these tests, we can conclude:

1. All three approaches allow to decrease the sampling rate if the SI is highly similar to the image to reconstruct.
2. Background subtraction is the most sensitive approach for the quality of the SI. Both the background subtraction approach and the $\ell_1\ell_1$ -minimization technique can have a negative impact if used with poor quality SI. This is not the case for CS with weighted SI.
3. When having multiple SIs with varying qualities, the CS approach using inter-SI and intra-SI weights will take advantage of the similarities while neglecting poor quality SI. This approach clearly outperforms the other techniques.

We can conclude that mm-wave imaging allows the Non-Destructive in-depth imaging of additive manufactured parts, using non-ionizing radiation in a stand-off setup. The high number of measurements can be reduced using a CS algorithm for the construction of the mm-wave images from the subsampled raw data. If SI is available, the SI can be exploited using a CS reconstruction with multiple weighted SI algorithm. The advantage of this approach over common techniques is that the quality of the SI does not need to be estimated in order to decide if the SI can be added without negative impact on the image quality. The algorithm autonomously neglects poor quality SI and exploits high quality SI at two levels: (1) between the multiple SI's and (2) between the elements of a single SI. The results obtained using as well the synthetic data as the real data, illustrate the important increase of the reconstruction quality compared to the common approaches.

Chapter 5

Through-the-Wall Imaging

5.1 Introduction

Through-the-Wall Radar Imaging is a research domain which gained a lot of interest over the last decennia, both within the scientific research community and the industry. This technology, which is in full development, is of interest for police, defense forces [Borek2005], fire and rescue personnel and first responders. Its primary goal is the detection, classification and tracking of humans and moving objects behind man made and visually opaque structures [Amin2017]. Recently, this type of radars has been used in Nepal, where four man, trapped under bricks, mud and debris after an earthquake, were found and rescued [Jet Propulsion Laboratory2015]. In a military and law enforcement context, this type of sensors is used in hostile rescue missions or to provide the soldier with situational awareness information before breaching a building in an urban environment [Farwell et al.2008]. Though the vast amount of research performed in this domain, many issues remain open and technical improvements are still needed in order to make this a mature and fully operational technology [Nkwari et al.2017]. In the remainder of this introduction we will first give a brief overview of the related research projects and main results in this area of research and subsequently we will discuss the state of the art of commercially available TWI sensors.

5.1.1 State of the art research on Through-the-Wall Imaging

We will now give an overview of the main research topics, prototypes of TWI sensors and results obtained on TWI published by various research institutes. We will confine our overview to the imaging of scenes hidden behind a wall and populated by humans. We will thus not treat the research performed on interior structure reconstruction [Ertin and Moses2009, Subotic et al.2008], where the aim is to build a blueprint of a building nor on vital body sign detection [Farwell et al.2008, Wang and Fathy2012] which focuses on the detection of heartbeat, respiration and movement.

The first results of TWI radars emerged in the late 1990s and grew from the field of Ground Penetrating Radars (GPR) [Amin and Ahmad2014]. The microwave region seemed to be the best suited part of the EM spectrum to accomplish the tasks of a TWI sensor due to the excellent penetration properties through man made structures [Frazier1997]. 2-D images, visualizing a horizontal slice through the illuminated room, were obtained with

sensors using ultrawideband pulses or Stepped Frequency Continuous Wave (SFCW) radars, combined with a large (synthetic) antenna array [Barrie2004, Thanh et al.2008]. Imaging of a scene behind a wall, assuming free space propagation revealed to be a naive approach since the interaction with the wall severely impacts the further propagation of the EM waves. Shadowing, attenuation, multipath, reflection, refraction, diffraction and dispersion effects must be taken into account in order to obtain high quality and reliable images [Amin and Ahmad2014]. We will now give an overview of the solutions proposed in the TWI radar literature. Table [5.1] summarizes the characteristics of the encountered TWI radar prototypes used in a research context.

Wall modelling

From the parameters describing the characteristics of the wall, such as the dimensions of the wall (length and thickness) and dielectric constant, the wall reflections can be modeled and subtracted from the received raw data. The work in [Dehmollaian and Sarabandi2008] presents a method for estimating the wall parameters from the first wave reflections. The drawback of this method is the need for a calibration step which must be performed under exactly the same conditions as the actual measurement. The work in [Leigsnering et al.2014] develops an analytical model for the received signal considering wall reverberation, multipath and signal attenuation.

Wall mitigation techniques

Wall mitigation techniques are developed to suppress the wall clutter dominating the radar image and are based on following assumptions: (1) the Round Trip Time (RTT) between the wall and the radar antenna is invariant over the antenna elements or antenna scanning positions in the case of a SAR, (2) the wall is homogeneous and large compared to the beam width of the antenna and (3) The wall is planar and causes specular reflections towards the radar antenna.

The work in [Yoon and Amin2009] applies spatial filters for mitigating the front wall reflection by removing the spatial zero-frequency and low-frequency components which are supposed to be caused by the wall reflections. Two filters are implemented, a Moving Average filter (MA) and a notch filter, and their performances are evaluated on real SFCW data. Both filters are reported to suppress the wall clutter significantly and thereby increase the Target-to-Clutter Ratio (TCR).

A different approach for removing the strong reflections of the front wall is proposed in [Tivive et al.2011a], by performing an eigen-analysis of the raw data through a singular value decomposition. Due to its strength, the wall reflections are generally present in the dominant singular vectors and their presence can be removed by projecting the measurements on the wall orthogonal subspace. This approach is further improved by the authors in [Tivive and Bouzerdoum2013] by suppressing the remaining wall clutter residual. Unfortunately, this technique was only tested on simulated data.

Multipath

Besides generating strong clutter, the walls are also at the origin of false positive detections. False targets, often referred to as ghosts, make their appearance on through-the-wall images and originate from multipath reflections between walls and targets. The work in [Ahmad2008], proposes an approach for determining if a target is real or must be neglected. The method requires radar images obtained from different locations along two sides of the room. The final image is then obtained through an RCS-based direct thresholding scheme for the fusion of the two images. [Leigsnering et al.2014] subdivides the multipath propagation into four categories:

1. Interior wall multipath: indirect paths originating from the reflection on one or more walls.
2. Floor and ceiling multipath: originating from secondary reflections on the floor or ceiling.
3. Wall ringing multipath: originating from multiple reflections within the front wall.
4. Interaction multipath: originating from target-to-target interaction.

and develops a model for the interior and wall ringing path for a scene with a single wall. The work in [Setlur et al.2011] and [Ma et al.2018] proposes a theoretical model for the interior wall multipath propagation in a room delimited by 4 perpendicular walls. This model is not only utilised to designate and mitigate the false targets, but also associates each ghost to its corresponding true target. The ghost signal is then mapped back to its true target which enhances the signal to clutter ratio of the target. Recently, a novel approach for ghost suppression and which is also immune to errors in the wall model was presented in [Guo et al.2018] and is based on array rotations. This method exploits the fact that the target location is independent of the array configuration while this is not the case for the location of the ghosts. Thanks to the array rotation, multiple images are generated and a ghost-free image is then obtained through incoherent multiplication fusion.

Compressed Sensing

The idea of applying Compressed Sensing on TWI measurements was first launched by Yoon and Amin in [Yoon and Amin2008] and was motivated by the observation that (I)SAR and TWI measurements require large bandwidth signals to be collected over a considerable amount of time or amount of antenna locations, resulting in a large number of data samples. Reducing the number of samples and at the same time obtaining images with the same image quality as obtained for fully sampled measurements, would thus reduce the measurement time. This reduction is important in TWI if only a short time window is available for performing the measurements or when the targets need to be stationary in order to avoid smearing and blurring. This work proves the potential of applying CS through a series of synthetic measurements but also stipulates that the technique was at that time not mature. The most important obstacle for CS to be effective for TWI is the fact that the image is not sparse, even if the number of targets is sparse, due to the wall effects described earlier in this introduction.

Subsequent to this first publication, the application of CS to TWI was further investigated in [Ahmad and Amin2012] and [Ahmad and Amin2013], where the sparsity of the image to reconstruct was obtained by using a change detection approach. The data used as measurement is obtained by the coherent subtraction of two measurements taken at a different moment in time. The work in [Leigsnering et al.2011] utilizes a 2D Discrete Wavelet Transform (DWT) as sparsifying bases. The ℓ_1 -norm of the DWT of the scene is minimized instead of imposing the sparsity of the through-the-wall image.

The work in [Lagunas et al.2012b] investigates if the wall clutter mitigation techniques maintain their performance when the measurements are undersampled and the image is reconstructed with a CS algorithm. In particular, the authors test how the SVD based and spatial filtering wall mitigation techniques perform. They report that those techniques do not lose in performance. However, it should be noted that the samples were taken uniformly over the bandwidth. Later in this chapter, we will prove that random sampling outperforms a uniform sampling strategy for radar imaging. If the frequencies are randomly chosen over the bandwidth, the above mentioned assumptions which forms the basis of the wall clutter mitigation no longer holds. The authors in [Ahmad et al.2013, Ahmad et al.2015] come to the same observation and propose a discrete prolate spheroidal sequence based wall clutter rejection scheme, which allows to capture the energy of the wall returns at each antenna position. This signal is then subtracted from the randomly sampled data in order to obtain an image of the targets [Ahmad et al.2013, Zhu and Wakin2015].

5.1.2 Commercial Through-the-Wall Imaging sensors

Multiple Commercial Of-The-Shelf (COTS) through-the-wall radar systems are available these days. An overview of the recent and today available systems are enlisted together with their specifications and output in tables 5.2 and 5.3. Some of them have been developed through a collaboration between a governmental organisation and a private company. This is the case for the Akela Standoff Through-wall Imaging Radar (ASTIR) which was developed by the company AKALA in close collaboration with the American National Institute of Justice. The image in table 5.2 shows the result of a measurement of a standing person walking inside a building. The system is capable of detecting stationary targets. Other systems, like the Prism 200, the ReTWis 5 and the Xaver 100 can only detect moving targets. The ReTWis 5 can perform the detection and measure the position in range and bearing of a human target by detecting the small Doppler frequency originating from the respiratory movements of the thorax. The Israeli company Camero, producer of the Xaver 100, developed recently two more advanced systems: the Xaver 400 and Xaver 800, both capable of detecting moving targets and plotting them together with stationary targets, like for example walls.

The most important characteristics shared by the different commercial of-the-shelf through-the-wall radars are:

1. The commercial systems are small and lightweight and do not use large antenna arrays or do not perform a scan along the wall. This illustrates that in order to be operational relevant, the system must be operable without the need to access every position along the wall and without long scanning times.

TABLE 5.1: Overview of TWI radar prototypes used in a research context

Ref	Radartype/ waveform	frequency	wall	targets	Processing method(s)
[Dehmollan and Sarabandi2008]	SFCW	1-2.5 GHz	concrete wall (19cm)	tri-hedral reflector	preprocessing: wall model subtraction; processing: SAR processing
[Yoon and Amin2009]	SFCW	1-3 GHz	hollow cin-der blocks; concrete wall; plywood-gypsum wall	dihedral	preprocessing: Spatial filtering; processing: Delay-and-Sum beamforming
[Tivive et al.2011a]	SFCW	2-3 GHz	plywood-air-drywall	dihedral tri-hedral cylinder	preprocessing: SVD; processing: Delay-and-Sum beamforming
[Peabody Jr et al.2012]	FMCW-MIMO	2-4 GHz	concrete wall (10-20cm) cin-der blocks	humans sphere	SAR processing
[Ahmad2008]	SFCW	0.7-3.1 GHz	wall of concrete blocks (14cm)	dihedral	image fusion of beamformed images
[Setlur et al.2011] [Lagunas et al.2012b] [Ahmad et al.2013] [Ahmad et al.2015]	SFCW	1-3 GHz	concrete wall (15cm)	metallic spheres	beamforming with a multipath model
[Leigsnering et al.2014]	SFCW	1-3 GHz	concrete blocks	aluminium pipe	multipath model and beamforming
[Ahmad and Amin2013]	pulsed	1.5-4.5 GHz	cement board (1cm)	moving person	CS with change detection
[Leigsnering et al.2011]	SFCW	0.7-3.1 GHz	concrete wall (14cm)	static wall target	CS with 2D DWI

2. The commercial systems use small bandwidths (except for the Xaver which uses frequencies going up to 10 GHz).
3. The commercial systems deliver range and detection information but do not deliver high-resolution radar images.
4. The commercial sensors are used to detect human targets or the movement of human targets.

We can thus conclude that there is a large gap between the TWI radar prototypes used in a research context and the recent commercial of-the-shelf through-the-wall radars. The research prototypes try to deliver high quality and high resolution images of metallic reflectors whereas the COTS systems deliver more rudimental images with range and detection information. The commercially available systems on the other hand are smaller (no large antennas or long scanning paths) and also uses small bandwidths.

Compressed Sensing can help to bridge this gap by reducing the number of samples and or number of antenna elements or antenna positions for a scanning sensor on the one hand and on the other hand delivering high-resolution images. But, to make the detection of human targets possible through most type of walls, CS will have to work in concert with efficient wall clutter and multipath mitigation techniques.

TABLE 5.2: Overview of commercial of-the-shelf TWI radars

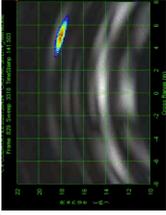
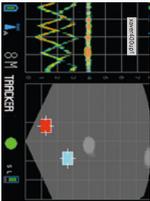
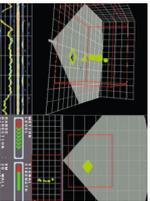
Product	Radartype/ waveform	frequency	walls	targets	output
ASTIR	SFCW	2.9-3.6 GHz	concrete walls	stationary and moving human targets	
ReTWis 5	SFCW	1.9-3.6 GHz	all building materials	detection of respiratory movements of humans	
Prism 200	FMCW	1.7-2.2 GHz	all types of walls	detection of moving objects/ humans	
Xaver 100	Pulse	3-10 GHz	all types of walls	detection of moving humans	

TABLE 5.3: Overview of commercial of-the-shelf TWI radars - continued

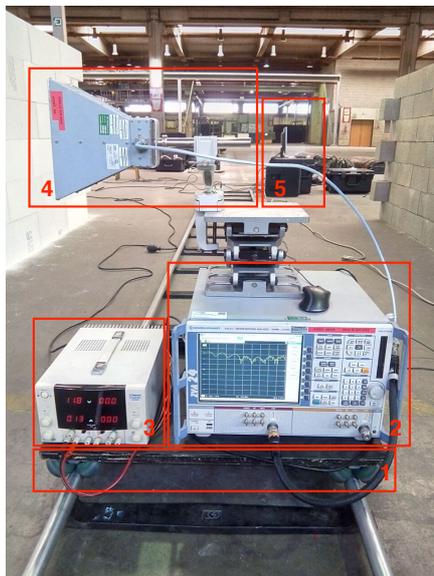
Product	Radartype/ waveform	frequency	walls	targets	output
Xaver 400	Pulse	2.9-3.6 GHz	all types of walls	detects moving humans and displays static objects targets	
Xaver 800	FMCW	3-10 GHz	all types of walls	detects moving humans and displays static objects targets	

5.2 Through-the-Wall radar experiments

In order to close the gap between the commercially available systems and the prototypes used in a research context, we will perform a series of experiments using Compressed Sensing on real TWI measurements. The ultimate goal is to reduce the number of samples that need to be taken in order to obtain a more practical system while preserving the image resolution, by reducing the measurement time. The real data is obtained through a series of tests using a TWI radar prototype that we built using a Vector Network Analyser (VNA). We start this section with a description of the prototype assembled for performing the TWI experiments, the choice of the parameters of our system and an explanation of the measurement strategy. We then describe the measurement setups, walls and scenes we built during our TWI radar measurement campaign and this section ends with the images obtained from these measurements through the application of a standard SAR compression algorithm.

5.2.1 Through-the-Wall radar prototype

Description



(a)



(b)

FIGURE 5.1: (a) A picture of the through-the-wall system setup with (1) the motorized board, (2) the VNA, (3) the power source, (4) the horn antenna and (5) the computer. (b) A screenshot of one of the windows of the user interface of the Labview program.

A through-the-wall SAR system was set up in the CISS department at the Royal Military Academy, for the experiments described and studied in this chapter. The self-built system is composed of:

- One motorized board and a rail
- One power source for the motorized board
- One Vector Network Analyzer (VNA): Rhode&Schwarz ZVA24
- One horn antenna: Schwarzbeck BBHA 9120
- One personal computer running Labview

Both the VNA and the motorized board are connected to the computer (via a GPIB and USB port respectively). We wrote a program with a graphical user interface (Figure 5.1 (b)) in Labview with following functionalities:

1. Input of the settings for the VNA:
 - start and stop frequencies and the number of equispaced frequencies within the chosen bandwidth
 - choice of the S-parameter: VNA port for the Input/Output
 - power
 - time gate: on/off, start and stop time
 - time gate type, type of filter, time gate shape and sideband suppression
2. Input of the parameters for the displacement of the motorized board:
 - velocity
 - total length of the displacement
 - number of measurement points
3. Launch the sequence for the SAR measurement and synchronize the displacements of the board and the measurements taken by the VNA.
4. Saving the raw data of all the measurements performed by the VNA over the whole scan.

The measurements are performed adhering to a stop and go method: the moving platform executes the displacement commanded by the Labview program along the rail and stops at a measurement position. Subsequently the VNA performs a measurement which is stored and the platform moves to the next sensor position to perform a next measurement. We have chosen to measure the S_{11} -parameter since this allows to use a single antenna for transmitting and receiving:

$$S_{11}(f_p) := \frac{S_{Rx}(f_p)}{S_E(f_p)} \quad [dB], \quad (5.1)$$

where $S_{Rx}(f_p)$ and $S_E(f_p)$ are the complex amplitudes of the received signal and the emitted signal at frequency f_p respectively. This choice has the advantage, compared to measuring the S_{21} -parameter, that we avoid misalignment problems and the crosstalk between the receiving and transmitting antennas.

The antenna is positioned with the smallest dimension horizontally. This choice is motivated by the fact that we want to obtain the largest opening angle in the plane we want to

image. Each point in the scene will thus be illuminated over a maximum number of positions. This way, we synthetically create the largest possible antenna and obtain a better resolution in the cross-range dimension.

TWI radar parameters

TWI radar parameter	Value
Starting frequency f_0	1 GHz
Bandwidth	4 GHz
Frequency step Δf	34.188 MHz
Scanning distance	4 m
Number of cross-range measurements	80
Number of frequencies	118
Aperture angle θ (-3 dB) at 3 GHz	40°
Range resolution	0.038 m
Cross-range resolution	0.078 m

TABLE 5.4

An overview of the parameters of the TWI SAR prototype is given in table 5.4. The choice of the utilized frequency band was made after measuring the transmission characteristics of the two types of material used to build the walls for the experiments. For this test, two antennas are placed facing each other, separated by 2 m and are connected to ports one and two of the VNA. We then measure the S_{12} transmission coefficient over the whole bandwidth of the antennas for three different configurations: (1) free space between the two antennas, (2) a wall of aerated concrete blocks placed perpendicular and in the middle of the two antennas and (3) a wall of hollow concrete blocks in between the two antennas. Figure 5.2 shows the outcome of these experiments. We observe first of all that the transmission coefficient decreases for increasing frequencies for the three experiments. This trend is most noticeable when the walls are in between the two antennas and we measure a serious reduction of the transmission through the blocks for frequencies above 5 GHz. This trend explains why most TWI radars operate at frequencies between 500 MHz and 5 GHz. We also observe a strong reduction of the transmission properties for the hollow concrete blocks between 2 and 3 GHz. This negative transmission peak disappeared whenever we changed the orientation of the hollow concrete blocks. This peak originates thus from the geometric properties of the blocks. After the evaluation of these experiments, we decided to perform the measurements between 1 GHz and 5 GHz.

The bandwidth of 4 GHz defines the resolution in the range dimension, which is equal to 3.75 cm. If we choose the size of the range bins equal to this resolution, i.e. we do not apply any oversampling, we need 118 range bins to cover a range equal to 4.425 m. If we want the linear set of equations, describing the measurements, to have a single solution for the reflectivity coefficient of each range bin, the number of samples needs to be equal to 118. We will thus select 118 equispaced frequencies over the bandwidth of 4 GHz. This explains the frequency step equal to 34.188 MHz.

The SAR TWI system scans the scene by performing a displacement in the cross-range dimension parallel to the wall over a total distance of 4 m. The cross-range resolution is calculated by:

$$\Delta r_{cross-range} = \frac{c}{2f \sin(\theta)} \quad (5.2)$$

and varies between 15.06 cm (for $f = 1$ GHz and $\theta = 95^\circ$) and 3.92 cm (for $f = 5$ GHz and $\theta = 50^\circ$). We approximate the obtained resolution in cross-range direction from the center frequency: 7.78 cm (for $f = 3$ GHz and $\theta = 40^\circ$). We decided to take measurements every 5 cm in order not to lose information along the cross-range dimension.

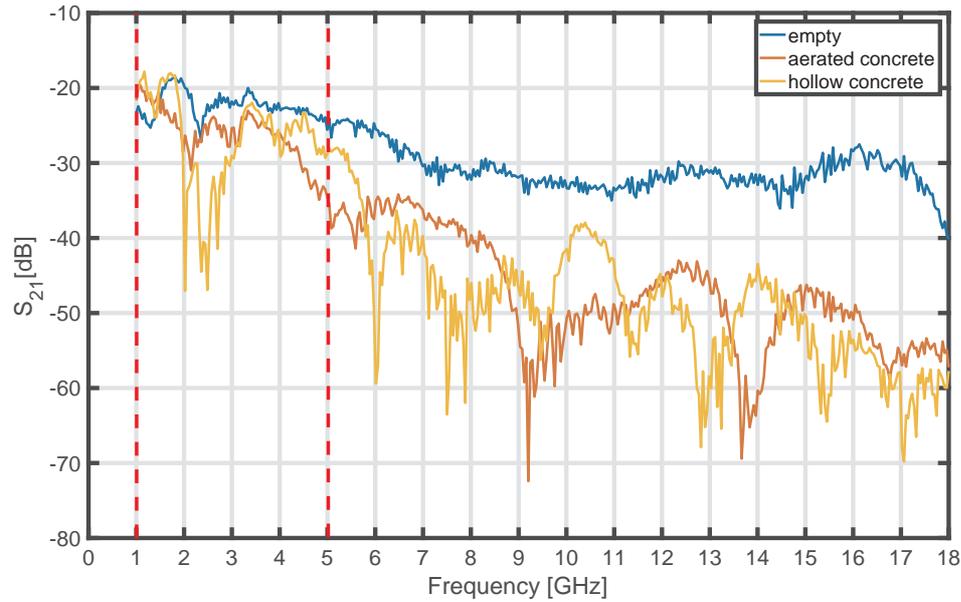


FIGURE 5.2: Transmission coefficient S_{21} in [dB] for free space, with an aerated concrete wall and with a hollow concrete wall between the two antennas.

5.2.2 Measurement setup

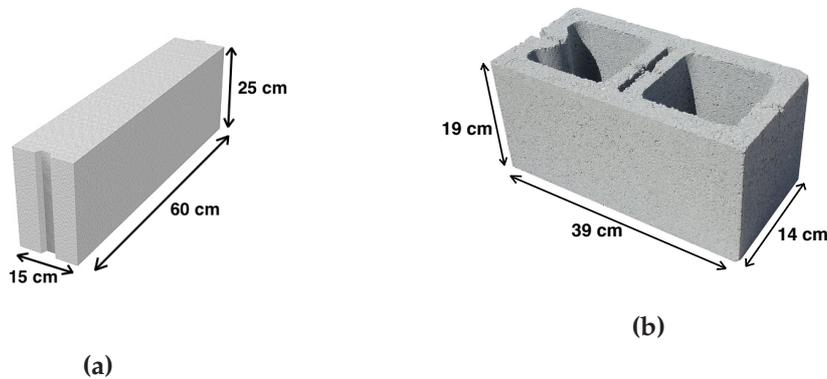


FIGURE 5.3: (a) Aerated concrete block and (b) Hollow concrete block used to build the walls for the TWI experiments.

We performed a series of tests with the TWI SAR prototype on scenes with self built walls. Two types of materials, used frequently for building external walls, were used : (1) Aerated concrete blocks (Figure 5.3 (a); depth: 15 cm, width: 60 cm and height: 25 cm) and (2) Hollow concrete blocks (Figure 5.3 (b); depth: 14 cm, width: 30 cm and height: 18 cm). The rooms built for the experiments are delimited by two walls parallel to the scanning track of the sensor. These walls were built by piling the described blocks without the use of mortar. This allowed us first to build, demolish and rebuild the walls with the same blocks and second we did not need to wait until the mortar dries before performing measurements. The rooms did not have walls in the across track dimension.

During the experiments the sensor was moving on a track parallel to the front wall and at a distance of 1 m. The walls are 4 m long and 1.5 m high. The back wall is at a distance of 3m from the front wall and parallel to the front wall. Figure 5.4 shows a photograph taken during one of the experiments. The sensed scene consists of a room with the presence of three human targets standing at different distances from the front wall. The data from this experiment will be used in the remainder of this chapter.

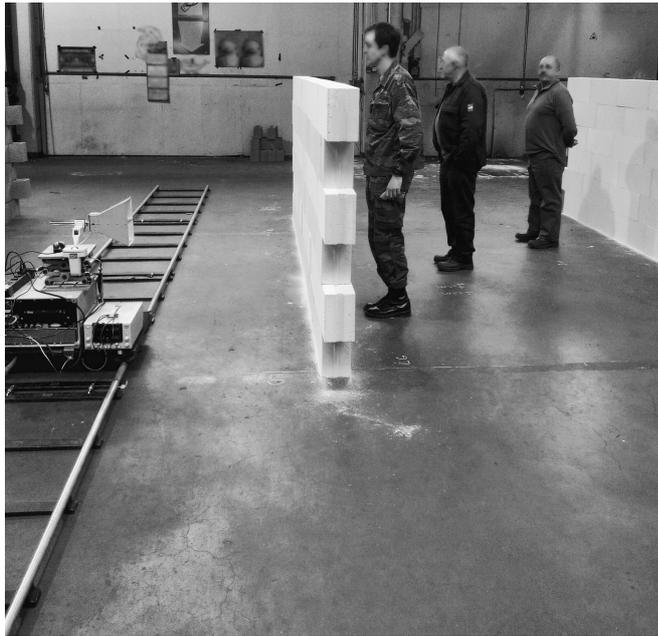


FIGURE 5.4: Example of an experiment with the TWI radar system.

5.2.3 Classic Synthetic Aperture Radar

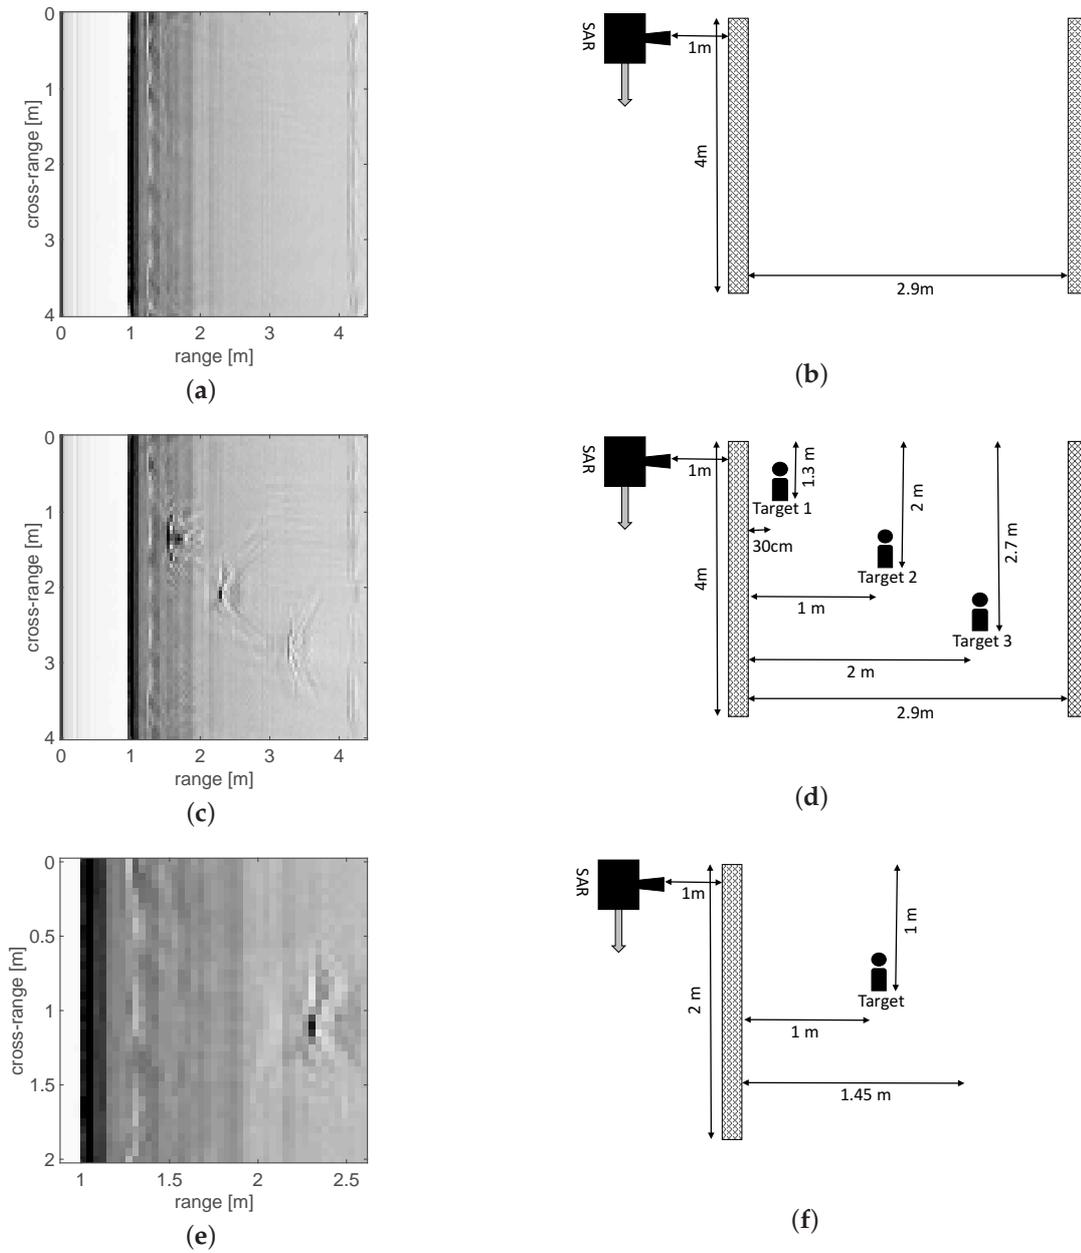


FIGURE 5.5: Images obtained with a Range Doppler SAR compression algorithm of three scenes together with a schematic representation of the scenes.

TWI experiments of three different scenes will be used in the next sections. Figure 5.5 depicts the SAR compressed images after applying a common Range Doppler compression

algorithm [Bamler1992] together with a schematic representation of the three scenes:

1. The first scene consists of an empty room delimited by two walls. This measurement will be used as background for performing coherent background subtraction.
2. A scene populated by three human targets standing at three different distances (30 cm, 1 m and 2 m) from the front wall. This experiment will be used throughout the remainder of this chapter.
3. A smaller scene of 2 m by 2.45 m with a single wall and a single human target at a distance of 1 m behind the wall. This smaller scene will be used for evaluating the performance of the wall mitigation techniques.

5.3 Compressed Sensing reconstruction

5.3.1 Through-the-Wall radar measurement matrix

The TWI radar measurement matrix is a SAR measurement matrix and is created as follows. For each resolution cell within the sensed scene, we create a raw data matrix populated by synthetic measurements of scenes where all resolution cells but one are characterised by a reflection coefficient equal to zero. These raw data matrices are then vectorised and form the columns of the measurement matrix. The number of columns corresponds thus to the number of resolution cells within the scene. The quality of the measurement matrix, which highly impacts the performance of the CS reconstruction, depends completely on the veracity of the synthetic measurements within the measurement matrix.

In chapter 3, each of the synthetic measurements performed for populating the measurement matrix, are obtained through a propagation model assuming free space conditions. The signal, after homodyne demodulation, coming from the reflection on a target (at position x), characterised by reflection coefficients a_l and received at sensor position (p), corresponds to:

$$s_p^x = \sum_{l=1}^L a_l \exp \left\{ -j2\pi f_l \frac{2R(x, p)}{c} \right\}, \quad (5.3)$$

where l is one of the L discrete frequencies emitted by the SFCW sensor and $R(x, p)$ is the euclidean distance between the antenna and the target. As explained in the introduction of this chapter, the free space propagation model is an inaccurate model once the signal reaches a wall. Figure 5.6 shows different unique round-trip paths between the sensor and a target [Leigsnering et al.2014]. In this section we will derive a more accurate model for the Round Trip Time (RTT) for EM waves passing through walls. The proposed model will take into account wall reverberations, transmission delays and attenuation but we will ignore interior wall and interaction multipath.

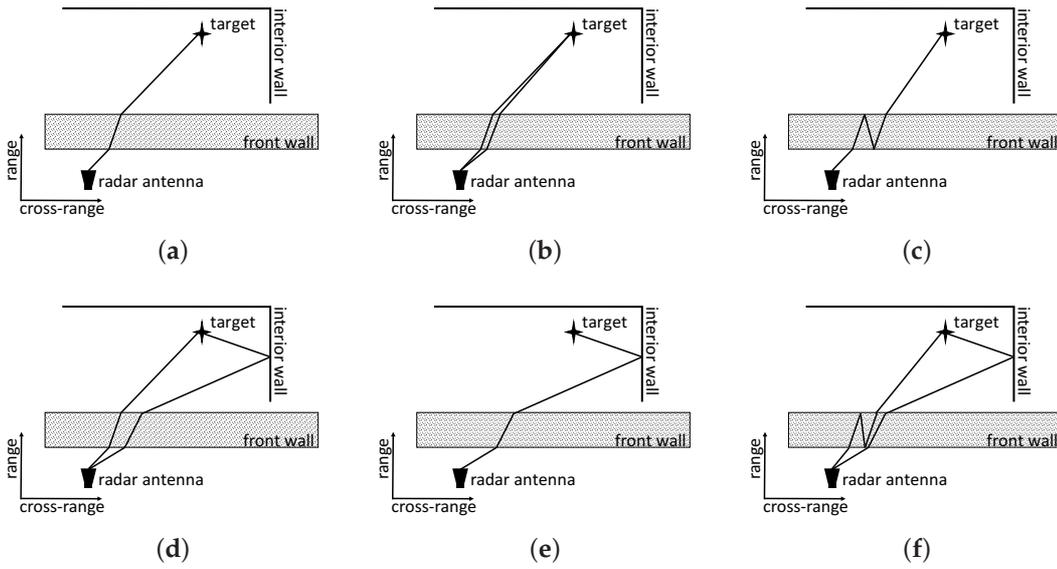


FIGURE 5.6: Round trip paths with maximum 1 reverberation and 1 multipath reflection.

Wall ringing multipath and propagation delay

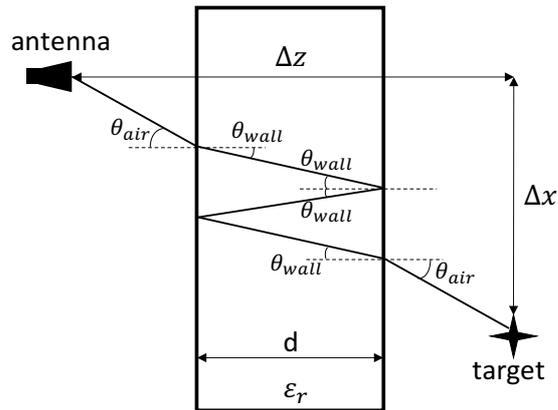


FIGURE 5.7: Schematic representation of the EM signal path through a wall.

Two wall parameters are supposed to be known for the derivation of the model:

1. ϵ_r : the relative permittivity ($= \epsilon / \epsilon_0$)
2. d : the thickness of the wall

First, an error on the location of the targets will be induced if not taking into account (1) the decreased velocity of the EM wave when propagating through the wall and (2) the refraction of the EM wave when the wave enters the front wall. Second, multiple reflections inside the wall (between the front and the back of the wall) will lead to the apparition of copies of the wall and or targets, known as wall ringing. Both phenomena will now be added to the model for the received reflections after propagation through a homogeneous, one layered wall.

The distance Δx in Figure 5.7 is equal to:

$$\Delta x = (\Delta z - d) \tan(\theta_{air}) + d(1 + 2i_w) \tan(\theta_{wall}), \quad (5.4)$$

where i_w is the number of internal reflections within the wall. The relationship between the angles θ_{air} and θ_{wall} is given by Snell's law:

$$\frac{\sin(\theta_{air})}{\sin(\theta_{wall})} = \sqrt{\epsilon_r}. \quad (5.5)$$

The nonlinear system of equations formed by the equations (5.4) and (5.5) can be solved using the Newton method. The one way propagation time becomes [Leigsnering et al.2014]:

$$\tau/2 = \frac{\Delta z - d}{c \cos(\theta_{air})} + \frac{\sqrt{\epsilon_r} d (1 + 2i_w)}{c \cos(\theta_{wall})}, \quad (5.6)$$

where for $i_w = 0$, we obtain the direct path.

Wall transmission coefficients

While propagating through the wall a considerable part of the signal will be absorbed and only a fraction will be transmitted. In [Abdeladi2015], the corresponding normalized transmission coefficients are deduced. The transmission coefficient for the direct path is equal to 1 and the coefficients for $i_w > 0$ are equal to:

$$\sigma = \frac{2 \cos(\theta_{air})}{\cos(\theta_{air}) + \sqrt{\epsilon_r} \cos(\theta_{wall})} \frac{2\sqrt{\epsilon_r} \cos(\theta_{wall})}{\cos(\theta_{air}) + \sqrt{\epsilon_r} \cos(\theta_{wall})} \left(\frac{-\cos(\theta_{air}) + \sqrt{\epsilon_r} \cos(\theta_{wall})}{\cos(\theta_{air}) + \sqrt{\epsilon_r} \cos(\theta_{wall})} \right)^{2i_w}. \quad (5.7)$$

The transmission coefficients for $i_w > 2$, tend towards zero and we will thus neglect the reflections after more than two reverberations inside the wall.

Antenna pattern

Antenna diagrams are defined in the case of TEM (Transverse Electromagnetic Mode) propagation, which means that the electric and magnetic fields are perpendicular and transverse to the direction of propagation of the EM wave. The area where this occurs is called the far field or Fraunhofer area. A sufficient condition for the minimum distance to the antenna is:

$$R \geq \frac{2D^2}{\lambda} \quad (5.8)$$

The distance is obtained at 1 GHz and is equal to 0.95 m.

The dimensions of the synthetic aperture created by the SAR system depend on the radiation pattern of the antenna. The larger the illuminated area, the larger the synthetic aperture and the better the cross-range resolution. The antenna pattern can be characterised by the -3 dB opening angle and we can consider to have a constant radiation within this whole angle and no radiation outside this -3 dB angle. This simple strategy for modelling the Antenna radiation pattern is naive, knowing that:

1. The -3 dB opening angle varies over the emitted bandwidth. We read for example in the technical specifications of the antenna an opening angle of 52.9° at 2 GHz and a much smaller angle of 32.4° at 4 GHz.
2. The radiated power over the different angles within the opening angle is not constant.

In order to obtain an accurate model for the behavior of the antenna, we performed a series of antenna measurements for measuring the antenna diagram. The antenna diagrams were obtained at the LEMA-laboratory for 201 frequencies over the total utilized bandwidth during the TWI experiments (1 to 5 GHz) and for angles ranging from -70° to $+70^\circ$. The normalized antenna diagrams obtained from these experiments at 1 GHz, 2 GHz, 3 GHz, 4 GHz and 5 GHz are depicted in Figure (5.8). We observe indeed an important change of the radiation pattern within the opening angle and over the different frequencies. We used a linear interpolation between two measured values for angles or frequencies for the antenna gain which we integrate in the equation (5.9).

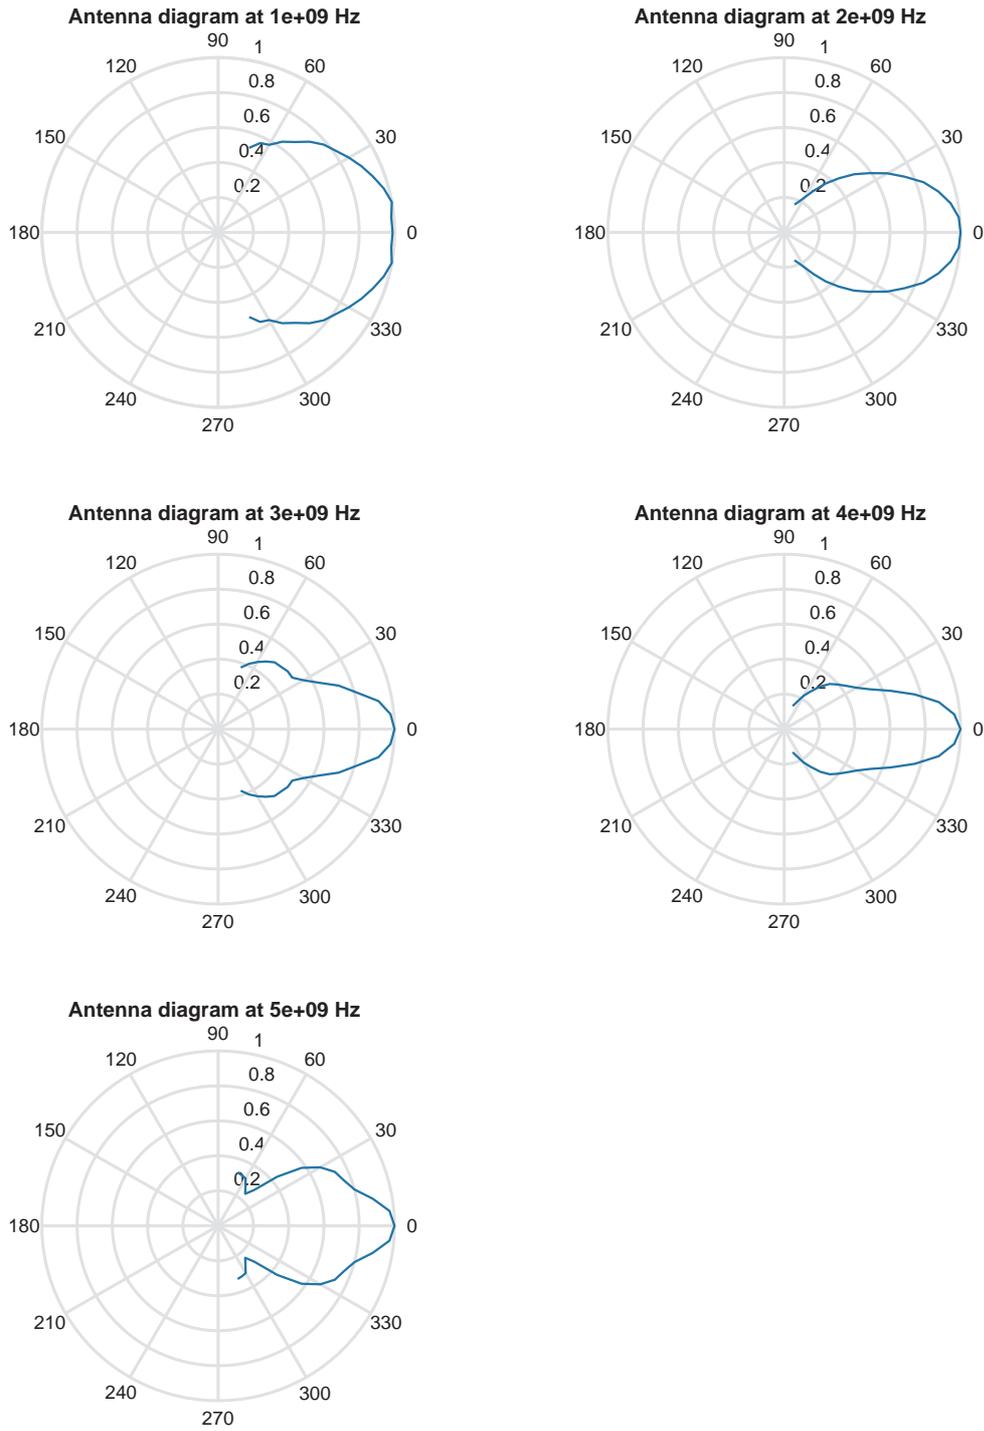


FIGURE 5.8: Antenna diagrams at 1 GHz to 5 GHz.

Received signal model

After adding the propagation model of the signals through the wall, the transmission coefficients and the antenna diagram to the received signal model, equation (5.3), becomes:

$$s_p^x = \sum_{i_w=0}^2 \sum_{l=1}^L c_{AE}^2(l, x, p) \sigma(i_w, \epsilon_r, x, p) \exp \{-j2\pi f_l RTT(i_w, \epsilon_r, d, x, p)\}, \quad (5.9)$$

where c_{AE} is the normalized antenna gain at frequency l , σ is the normalized wall transmission coefficient calculated from (5.7) and the τ is obtained from (5.6).

5.3.2 Compressed sensing reconstruction

We now reconstruct the TWI SAR image obtained for the scene 5.5 (d), using a Compressed Sensing algorithm. We do not exploit the structure in the sparsity. One can easily observe that the non-sparse components in the images are clustered. This additional structure could be exploited by using CS for block sparse signals [Eldar et al.2010]. The measurement matrix A is obtained using the signal model described by equation (5.9). The CS algorithm used is *l1qc-logbarrier* from the package ℓ_1 -magic, a log-barrier algorithm for second-order cone programs. The results obtained using a decreasing number of randomly selected frequencies, is shown in Figure 5.9. We observe visually that the image quality rapidly decreases for decreasing subsampling rates. With 70% of the original number of samples, the targets are already drown in the noise. We remark that when the subsampling rates decrease, the almost constant background clutter level around the targets becomes salt and pepper noise with very dark and very light pixels.

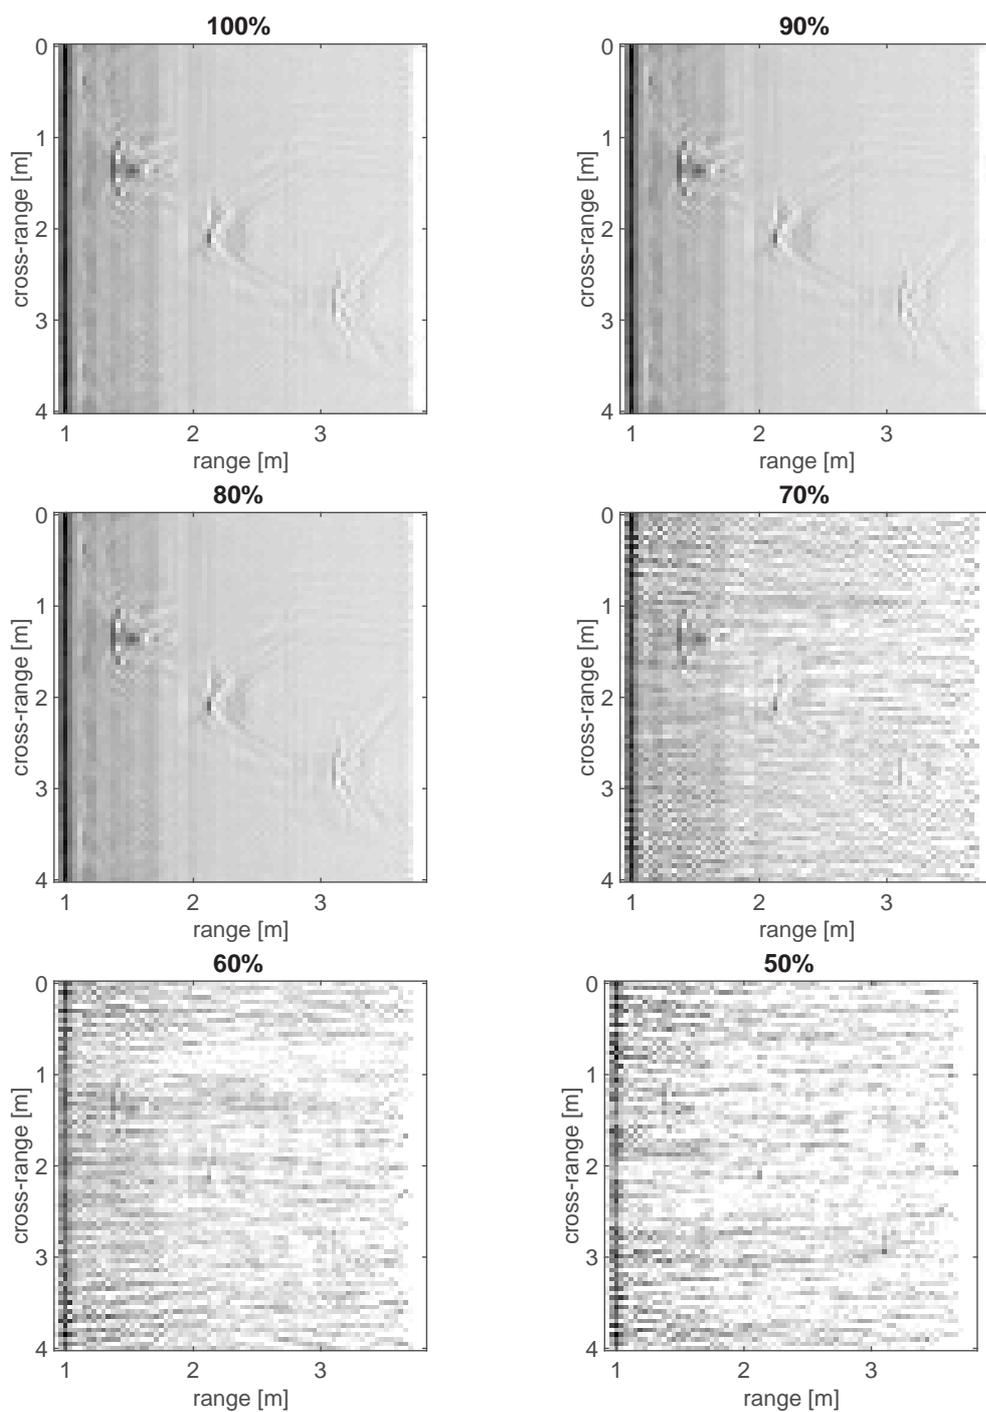


FIGURE 5.9: CS reconstruction (using ℓ_1 -magic) for undersampling rates ranging from 100% down to 50%.

Evaluation of the reconstruction performance

We now want to quantify the image reconstruction performance of the CS reconstruction for varying subsampling rates. Popular metrics used for this task are:

1. The probability of detection for a fixed false alarm rate. CA-CFAR algorithms have been applied in [Khamlichi et al.2012, Debes et al.2010, Sun et al.2014] for this purpose. We implemented a CA-CFAR algorithm as follows: First, we select the pixel with the maximum value of each target. Second, we determine the size of the guard zone which will not be used to calculate the detection threshold. Finally we choose the dimensions of the zone with training cells which will be used to determine the detection threshold, supposing Gaussian noise:

$$Threshold = N_T(P_{fa}^{-1/N_T} - 1)P_n \tag{5.10}$$

with N_T the number of training cells, P_{fa} is the probability of false alarm and P_n is the average pixel values over the training cells. When performing this test on the scene reconstructed with 100% of the samples, the three targets are not detected, even for low probabilities of false alarm. The hypothesis of Gaussian noise is probably not valid.

2. Target to Clutter Ratio. TCR was used as a metric for the evaluation of TWI in [Lim and Nam2014, Narayanan et al.2017, Lagunas et al.2012a] and is calculated as:

$$TCR = 20 \log_{10} \left(\frac{\max_{(k,l) \in A_t} |b(k,l)|}{\frac{1}{N_c} \sum_{(k,l) \in A_c} |b(k,l)|} \right), \tag{5.11}$$

with A_t the target area, A_c the clutter area, N_c the number of elements within the clutter area and $b(k,l)$ the pixel value in element (k,l) .

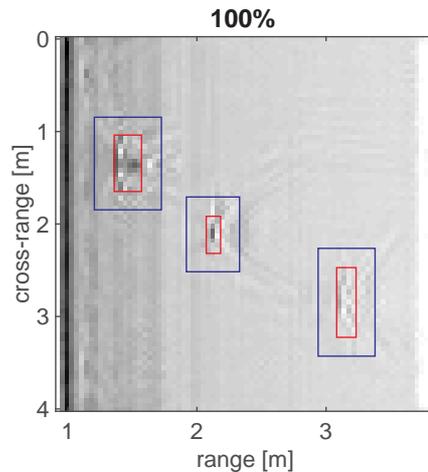


FIGURE 5.10: SAR image of the scene with framed in red: the target areas and in blue: the clutter areas.

Figure 5.10 shows the chosen target areas (within the red rectangle) and clutter areas (within the blue rectangle). The corresponding TCRs are listed in table 5.5 together with the maximum value within the target area and the mean of the clutter

TABLE 5.5: TCR for the three targets.

	Target 1		Target 2		Target 3	
	100%	50%	100%	50%	100%	50%
TCR [dB]	6.90	10.85	9.62	13.97	7.37	18.39
Max(A_t)	4.8E-3	4.9E-3	4.4E-3	4.0E-3	2.7E-3	5.3E-3
Mean(A_c)	2.1E-3	1.4E-3	1.5E-3	7.9E-4	1.1E-3	6.4E-4

area. Although, it is visually unequivocally clear that the targets are no longer visible at a subsampling rate of 50%, the TCR at 50% is higher than the TCR obtained at 100% for the three targets. We observe that the targets disappear in the spiky clutter, which explains the high maximum in the target area and the low mean value in the clutter area. For our case, the TCR can thus not be used as a reliable metric to measure the reconstruction performance.

3. Peak Signal-to-Noise Ratio (PSNR) was used as a metric for evaluating CS applied on TWI in [Gaikwad and Shevada2013] and [Mohsin Riaz and Ghafoor2012]. The reference image used to calculate the PSNR is the image obtained with 100% of the samples and serves as a pseudo-truth. Figure 5.11 shows the mean PSNR of 32 reconstructions of the scene with three human targets obtained with an increasing number of samples going from 5% of samples up to 90% of samples. We notice that the PSNR strongly decreases between 65% and 75%, which corresponds with what we observe visually in Figure 5.9.

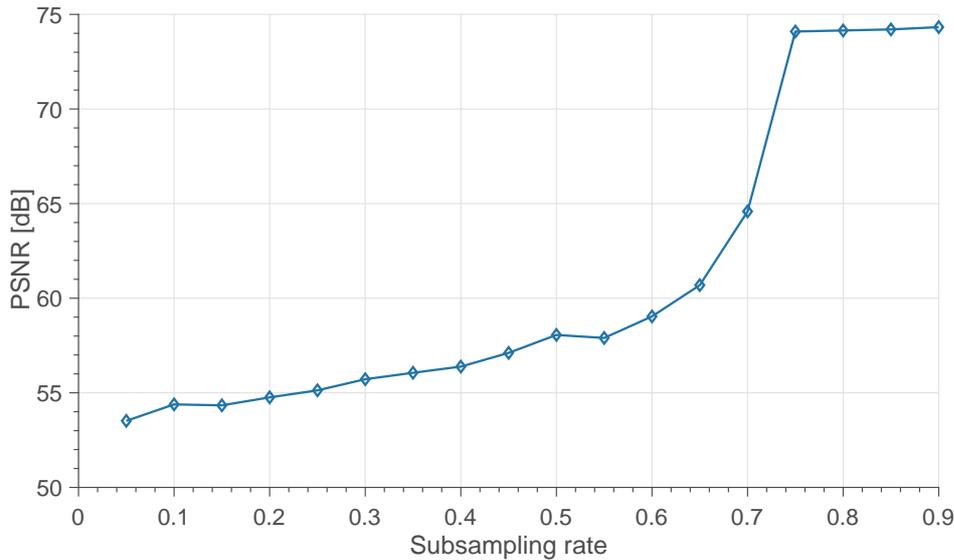


FIGURE 5.11: PSNR of the CS reconstructed TWI measurement as a function of the subsampling ratio

5.3.3 Sampling strategies

The results presented in the previous subsection are obtained when applying a random selection of frequency samples over all the emitted frequencies and all the sensor positions. When applied in a real environment, some operational constraints like the accessibility of the sensor positions, the availability of frequencies due to for example jamming or authorizations, can dictate a different subsampling strategy. We will now evaluate the impact of the sampling scheme on the reconstruction performance. We consider 4 sampling schemes:

1. Random frequency sampling in range and cross-range dimension (Figure 5.12 (a)).
2. Random sensor position sampling. The raw data contains the full received signal sampled at the Nyquist rate from a random number of sensor positions (Figure 5.12 (b)).
3. Random frequency sampling which is fixed over the different sensor positions. At each position, the reflectivity of the scene is measured at the same random frequencies (Figure 5.12 (c)).
4. A combination of sampling strategies (1) and (2) (Figure 5.12 (d)).

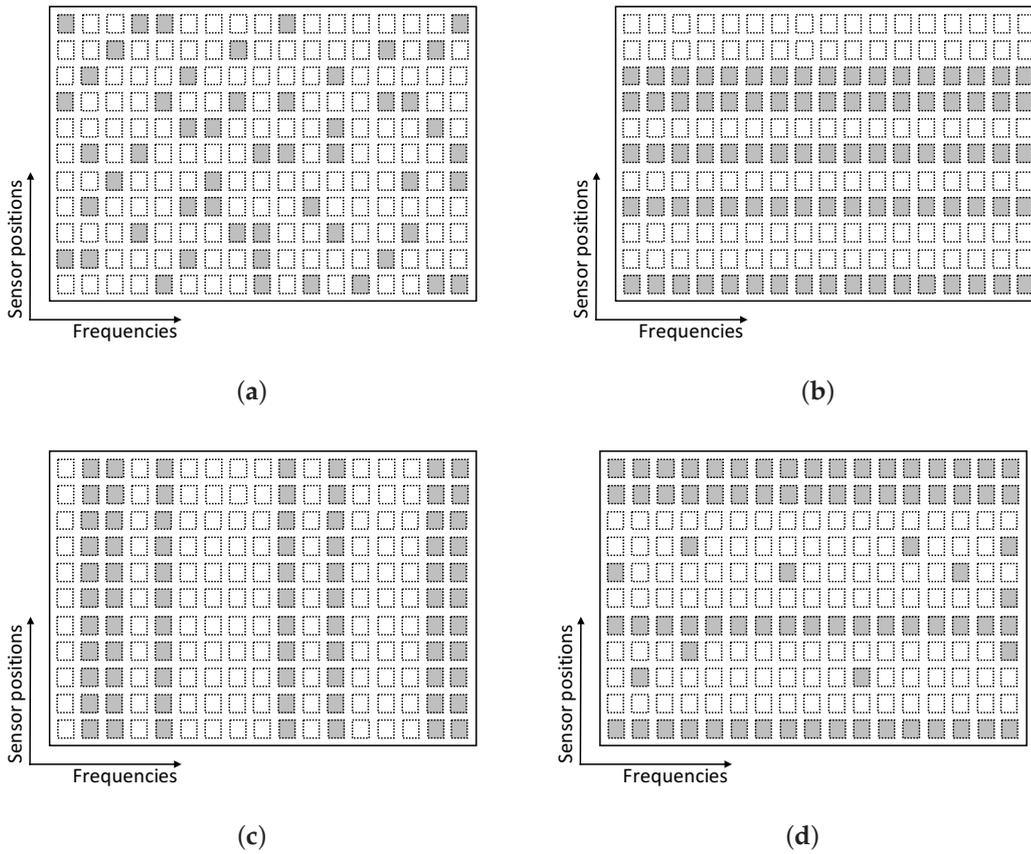


FIGURE 5.12: Sampling strategies

The performance of the CS reconstruction of the TWI scene with the three human targets, applying the four sampling strategies is measured as the PSNR and is depicted in Figure 5.13. The best result is obtained when adopting a random frequency sampling strategy in range and in cross-range dimensions. The worst result is obtained for using the same frequencies over all the sensor positions.

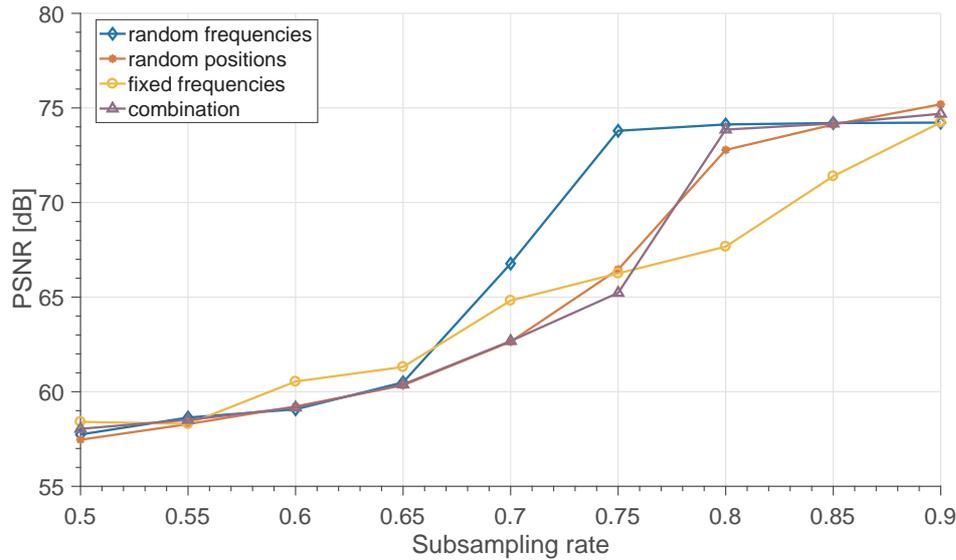


FIGURE 5.13: PSNR of the CS reconstructed TWI obtained with four different sampling strategies.

5.3.4 Sparsifying bases

The mediocre results obtained with the CS reconstruction of the TWI scene for decreasing subsampling rates can be ascribed to the fact that the TWI is not sparse. We will explore different solutions to circumvent this problem. A general solution for using CS for reconstructing non-sparse signals is the application of a sparsifying transform on the signal to reconstruct. Instead of reconstructing the non-sparse signal, we reconstruct the sparse transformation of that signal. After the CS reconstruction, the signal is obtained by applying the inverse transform.

We evaluate the use of two different sparsifying transforms:

1. The Daubechies-1 wavelet transform or Haar wavelet.
2. The Discrete Cosine Transform (DCT).

If we reconstruct the transformed signal, the measurement matrix needs to be right multiplied with the inverse transform. Instead of creating the transformation matrices explicitly, we use the Matlab function 'handles' for creating implicit matrices, which are accepted as inputs by the ℓ_1 -magic package. This way, we are able to use the Matlab transformation functions and do not have to write transform and inverse transform matrices.

Figure 5.14 depicts the histograms of the normalized pixel values of (a) the original TWI scene, (b) the wavelet transform and (c) the DCT transform of the same scene. We see

that the sparsity (the number of almost zero valued pixels) is increased for the wavelet transform and is even more increased after applying the DCT transform.

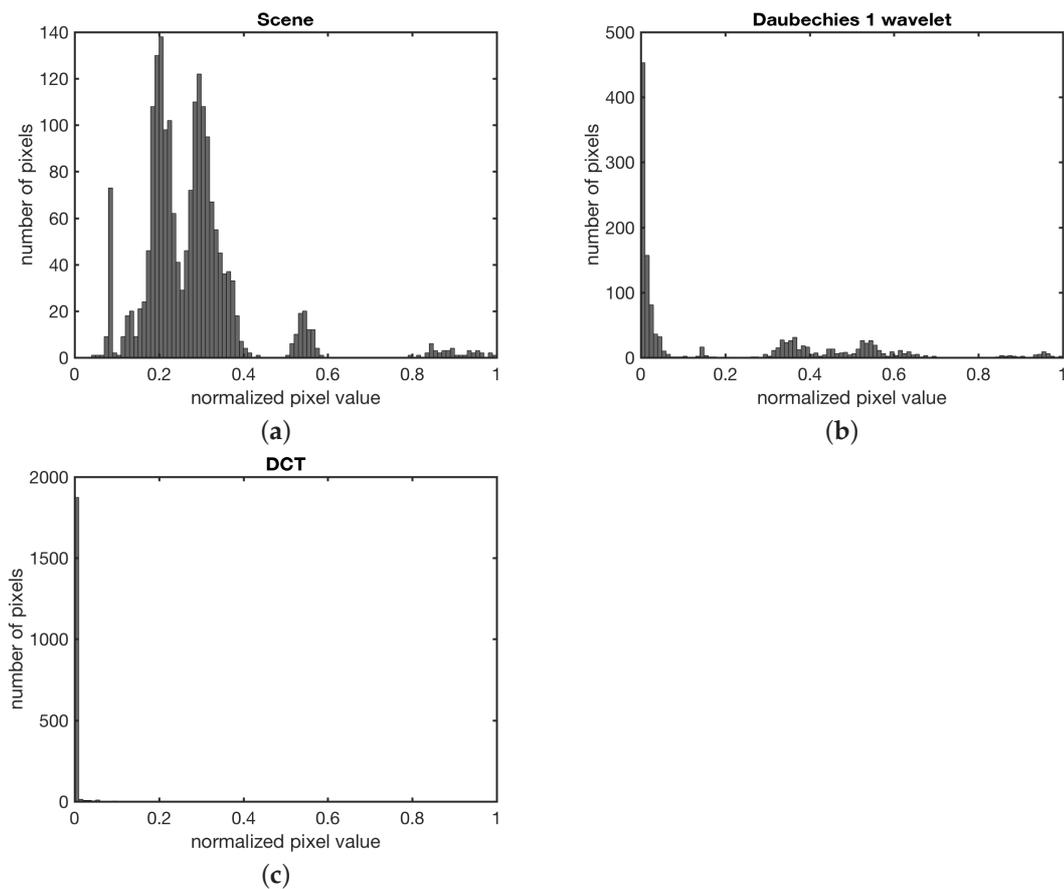
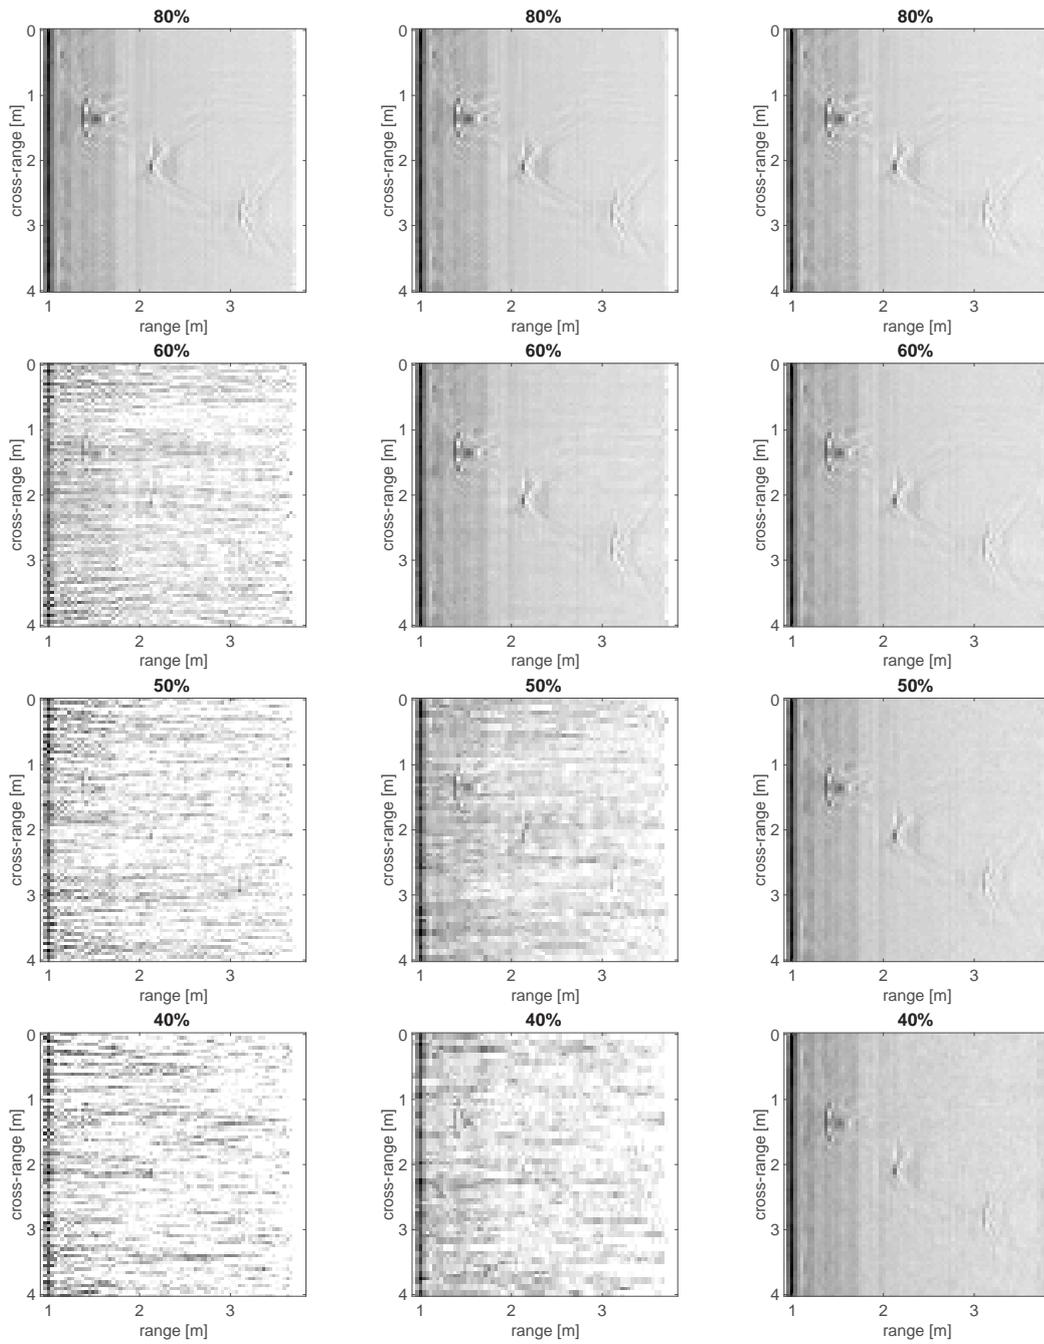


FIGURE 5.14: Histograms of the normalized pixel values of (a) the original TWI scene, (b) the wavelet transform of the scene and (c) the DCT of the scene.



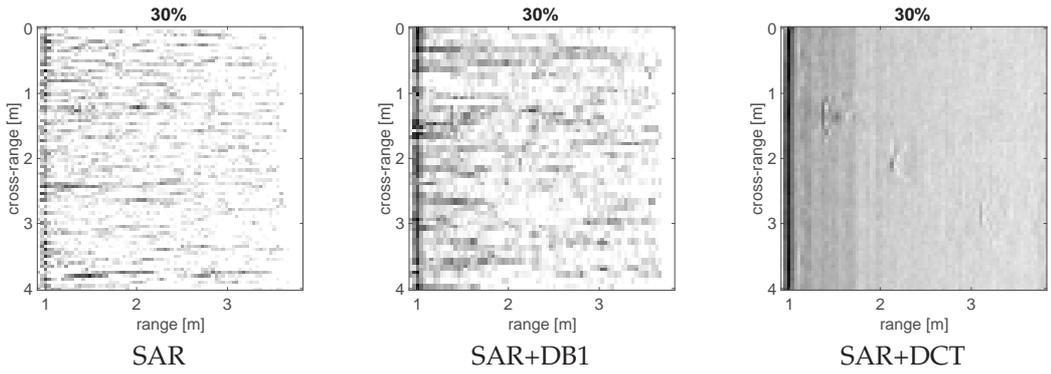


FIGURE 5.16: CS reconstruction, left column: without transform, middle column: using a wavelet transform and right column: using the DCT.

Figure 5.16 depicts the resulting images obtained without the use of a sparsifying transform (left column), by applying the wavelet transform (middle column) and by using the DCT (right column), for a subsampling rate equal to 80% and decreasing down to 30%. We observe that the targets remain visible for a subsampling rate of 60% in the case of the wavelet transform and start to fade away only around 30% when applying the DCT sparsifying transform. These conclusions are confirmed in Figure 5.17, showing the PSNR for increasing subsampling rates. Whereas the PSNR starts to decrease drastically for subsampling rates beneath 75% when not applying a sparsifying transform, the PSNR stays almost constant down to subsampling rates equal to 60% for the wavelet transform. When applying a DCT, the PSNR remains higher down to a subsampling rate of 25%.

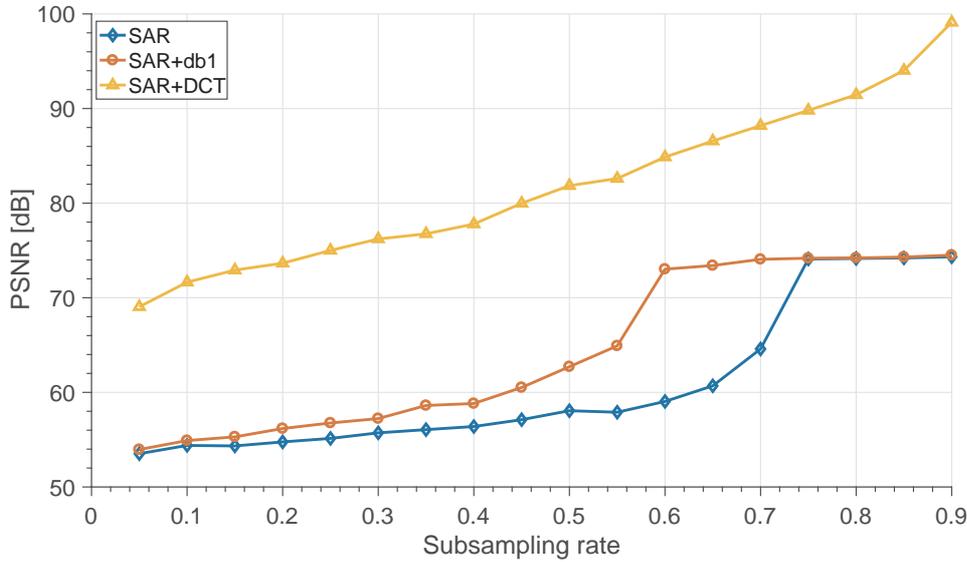


FIGURE 5.17: PSNR for increasing subsampling rates without using a sparsifying transform, with the wavelet transform and with the DCT.

5.3.5 Wall mitigation techniques

If the following assumptions can be considered to be true: (1) The front wall is parallel to the scanning path, (2) The front wall is large compared to the dimensions of the scene, (3) The front wall is homogeneous and the front wall reflections dominate the rest of the reflected signal, several techniques can be utilized in order to suppress the wall reflections and thereby mitigate the wall clutter. We will now evaluate four different techniques: (1) Coherent background subtraction, (2) the Moving Average filter (MA); (3) Singular Value Decomposition (SVD) and (4) Robust Principal Component analysis (RPCA). We will evaluate these techniques when reconstructing the fully sampled data and we will also evaluate the robustness of the techniques against subsampling. The performance will be evaluated by calculating the Target to Clutter Ratio (TCR). The TCR reveals to be a good metric for the wall mitigation techniques when using the fully sampled data, which is not the case for subsampled data as explained before. We will evaluate the resulting image visually when applying a random subsampling of the raw data.

The data used for evaluating the wall mitigation techniques are the real TWI measurements of the scene containing a single wall and a single target at a distance of 1 m behind the wall (Figure 5.5 (c)). Figure 5.18 shows the reconstructed images of this scene for sampling rates equal to 100%, 80%, 70% and 60% without using one of the mitigation techniques. We see again that the target starts to vanish for subsampling rates below 80%.

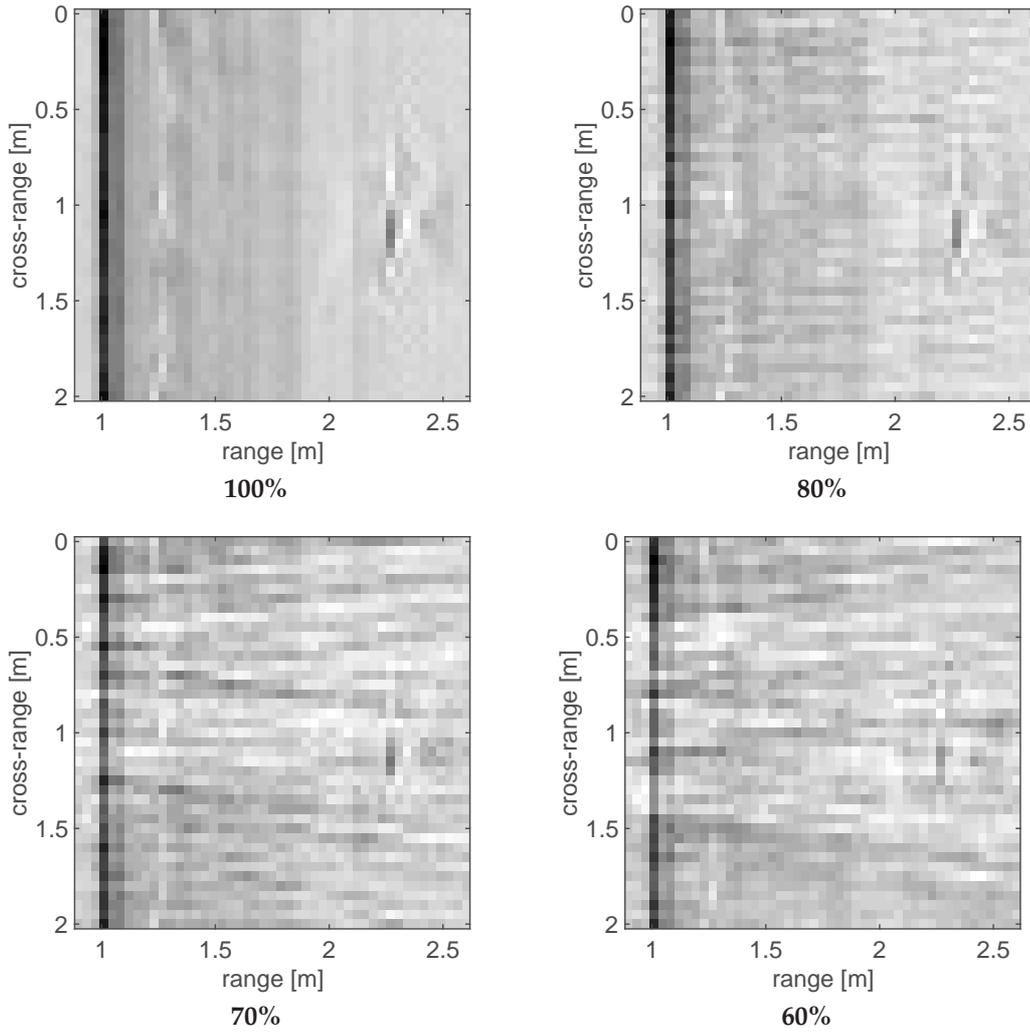


FIGURE 5.18: Reconstruction of the TWI scene with a single target for subsampling rates equal to 100%, 80%, 70% and 60%.

(1) Coherent background subtraction:

Coherent background subtraction is a very simple but effective approach to suppress the wall clutter. Its biggest drawback is the need for a measurement performed at a different moment in time (when the target was not at the same spot in the scene) but under the exact same measurement conditions. The background measurement used for this experiment was performed right after the actual measurement of a room left completely empty. The background measurement is then coherently subtracted from the measurement of the scene with the target. The CS reconstruction is applied on the measurement after the subtraction. The obtained image contains thus only the impact of adding the target.

When comparing the images obtained from the fully sampled data, the TCR increases from 5.73 dB for the original image towards 32.93 dB after applying coherent background subtraction. We also observe that the technique is very robust against subsampling. At a subsampling rate as low as 10%, the target remains clearly visible (Figure 5.19).

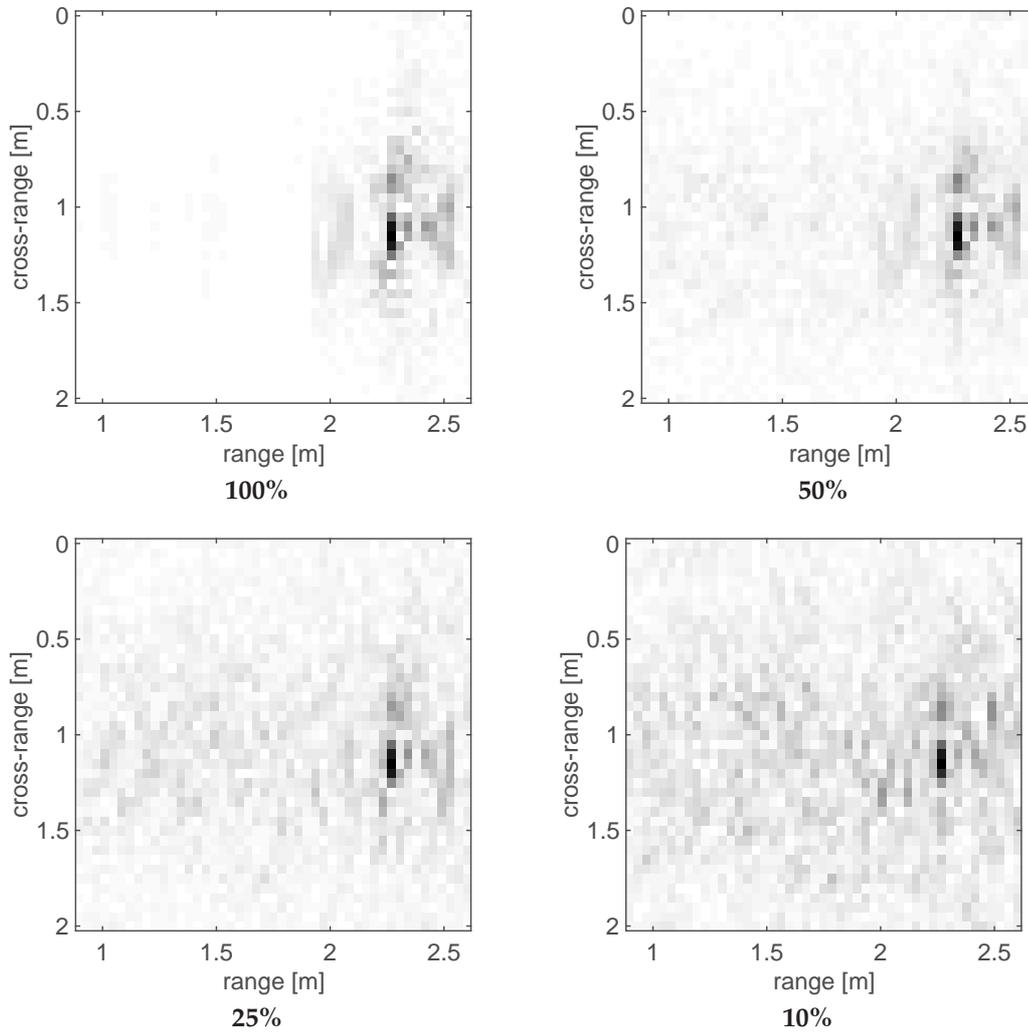


FIGURE 5.19: CS Reconstruction with coherent background subtraction for sampling rates equal to 100%, 50%, 25% and 10%.

(2) Moving average filter:

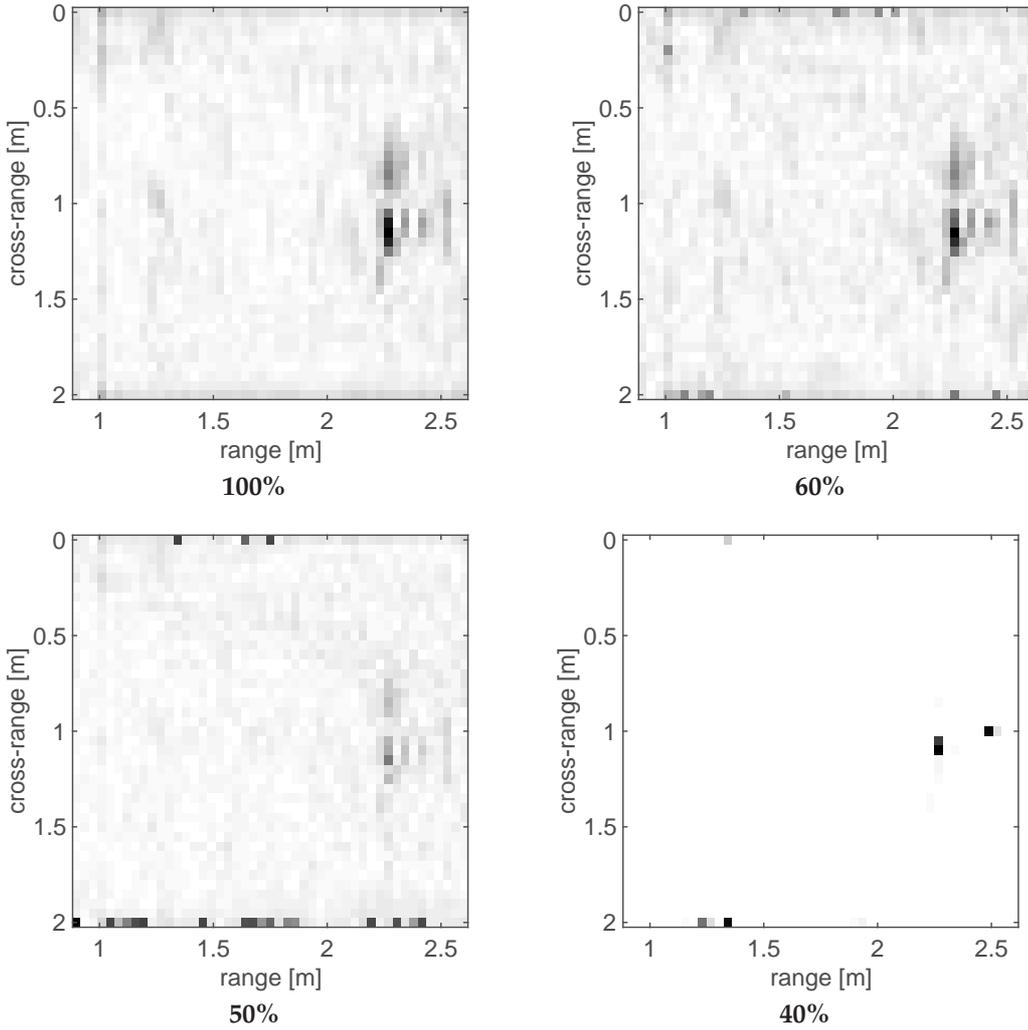


FIGURE 5.20: Reconstruction of the TWI scene with a single target, after applying a moving average filter on the raw data, for subsampling rates equal to 100%, 60%, 50% and 40%.

The total signal received by the radar at the scanning position p can be decomposed into two signals received after the reflection on the wall or on a target:

$$s(p, t) = s_w(t - RRT_{wall}) + s_t(t - RTT_{target}), \quad (5.12)$$

where s_w is the signal, described by equation (5.9), reflected by the front wall in the direction of the sensor and received after the Round Trip Time (RRT_{wall}) between the sensor and the wall and s_t is the signal, described by equation (5.9) and received after reflection

on a target. Under the previous made hypothesis, the reflection coming from the wall can be supposed to be invariant over the different sensor positions and can thus be filtered out by using a moving average filter, which is a simple non-causal Finite Impulse Response (FIR) filter, used here to filter out the components with a spatial frequency equal to zero.

The signal at frequency f_q and at sensor position p , after applying the MA filter is:

$$z_{MA}(p, f_q) = z(p, f_q) - \bar{z}(f_q), \quad (5.13)$$

where:

$$\bar{z}(f_q) = \frac{\sum_{m=1}^N z(m, f_q)}{N}, \quad (5.14)$$

if taking the whole scanning path into consideration. The work in [Amin and Ahmad2014] demonstrates that this corresponds to removing the spatial frequency component $\kappa = 0$.

Taking the discrete Fourier transform $Z_{MA}(\kappa, f_q)$ of (5.13):

$$Z_{MA}(\kappa, f_q) = \sum_{m=1}^N \left(z(m, f_q) - \bar{z}(f_q) \right) \exp\{-j2\pi\kappa m / N\}, \quad (5.15)$$

we obtain that:

$$Z_{MA}(\kappa, f_q) = \begin{cases} Z(\kappa, f_q) & \text{if } \kappa \neq 0 \\ 0 & \text{if } \kappa = 0 \end{cases} \quad (5.16)$$

The images in Figure 5.20 show the results obtained after applying a MA filter on the raw data for subsampling rates equal to 100%, 60%, 50% and 40%. At a subsampling rate equal to 100% we measure an increase of the TCR from 5.73 dB to 24.89 dB thanks to the application of a moving average filter. We further observe that the target is not longer visible for sampling rates below 60% despite the sparsity of the image after applying a MA filter. This can be explained by the random sampling of the frequencies over the cross-range dimension. By selecting different frequencies over the sensor positions, the invariant characteristic in the cross-range dimension starts to disappear.

(3) Singular value decomposition:

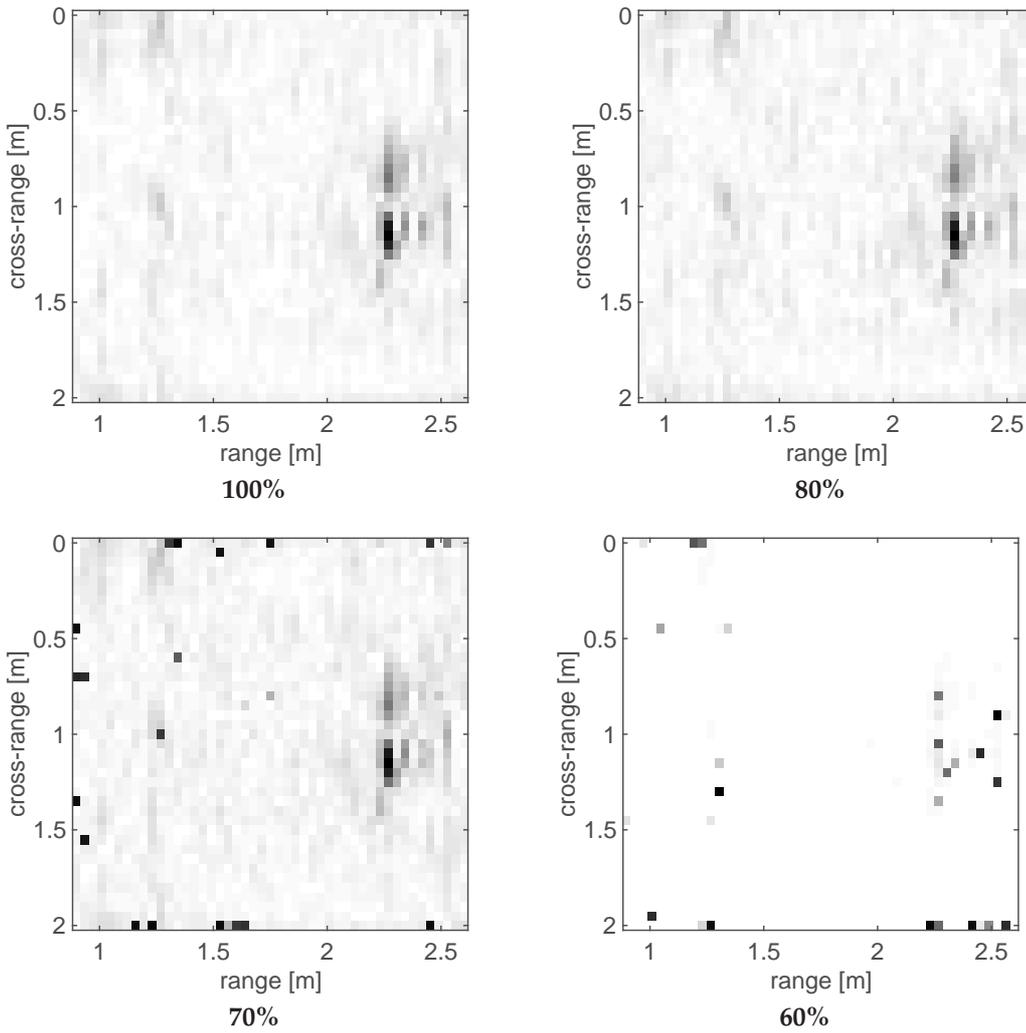


FIGURE 5.21: Reconstruction of the TWI scene with a single target using SVD for subsampling rates equal to 100%, 80%, 70% and 60%.

Singular Value Decomposition is a common technique used in signal processing to suppress the noise by performing a SVD of the data and subsequently retain only the dominant singular values. We will now apply SVD in the opposite way, by eliminating the dominant singular value(s) we try to remove the high clutter originating from the front wall reflection which is considered to be the strongest reflection.

Figure 5.21 depicts the obtained images after applying an SVD on the raw measurement data and eliminating the most dominant singular value. When removing also the second

most dominant singular value, the target immediately disappears. The TCR obtained from the 100% sampled measurement is equal to 26.69 dB. But, the reconstruction performance rapidly decreases for decreasing subsampling rates.

(4) Robust Principal Component Analysis:

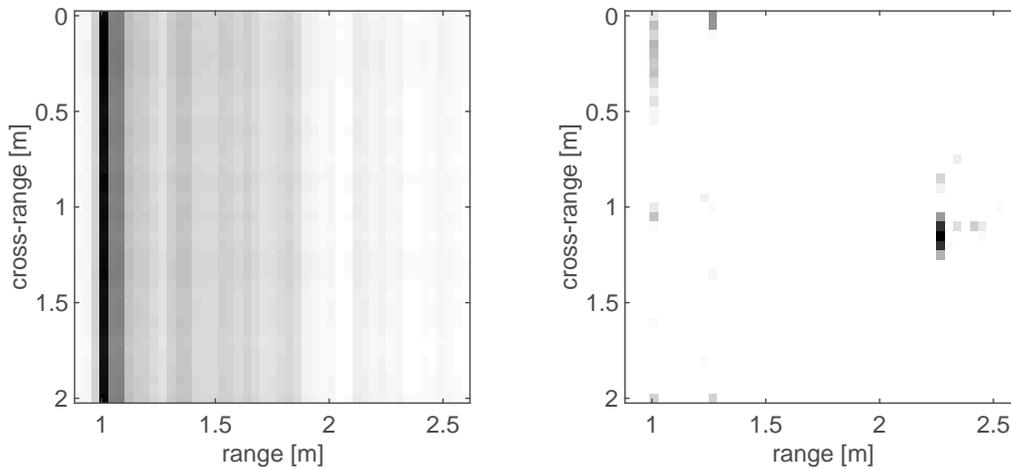


FIGURE 5.22: RPCA reconstruction of the TWI measurement with on the left: the background and on the right: the foreground with the target.

In contrast to the previously discussed wall mitigation approaches, where a mitigation technique is applied on the raw data followed by a regular CS reconstruction, the Robust Principal Component Analysis (RPCA) performs both at the same time. The application of RPCA results in two images: (1) a low rank background image and (2) a sparse foreground image. Figure 5.22 shows the obtained foreground and background images after applying RPCA on the fully sampled measurement. In contrary to other wall suppression techniques, we exploit the invariance in the image instead of in the measurement along the across-range dimension which implies that this technique is more robust to subsampling compared to the other wall mitigation approaches.

Figure 5.23 shows the reconstructed images using RPCA. We observe that the target remains clearly visible down to a subsampling rate equal to 30%. The TCR of the foreground image obtained using 100% of the original number of samples is equal to 52.65 dB.

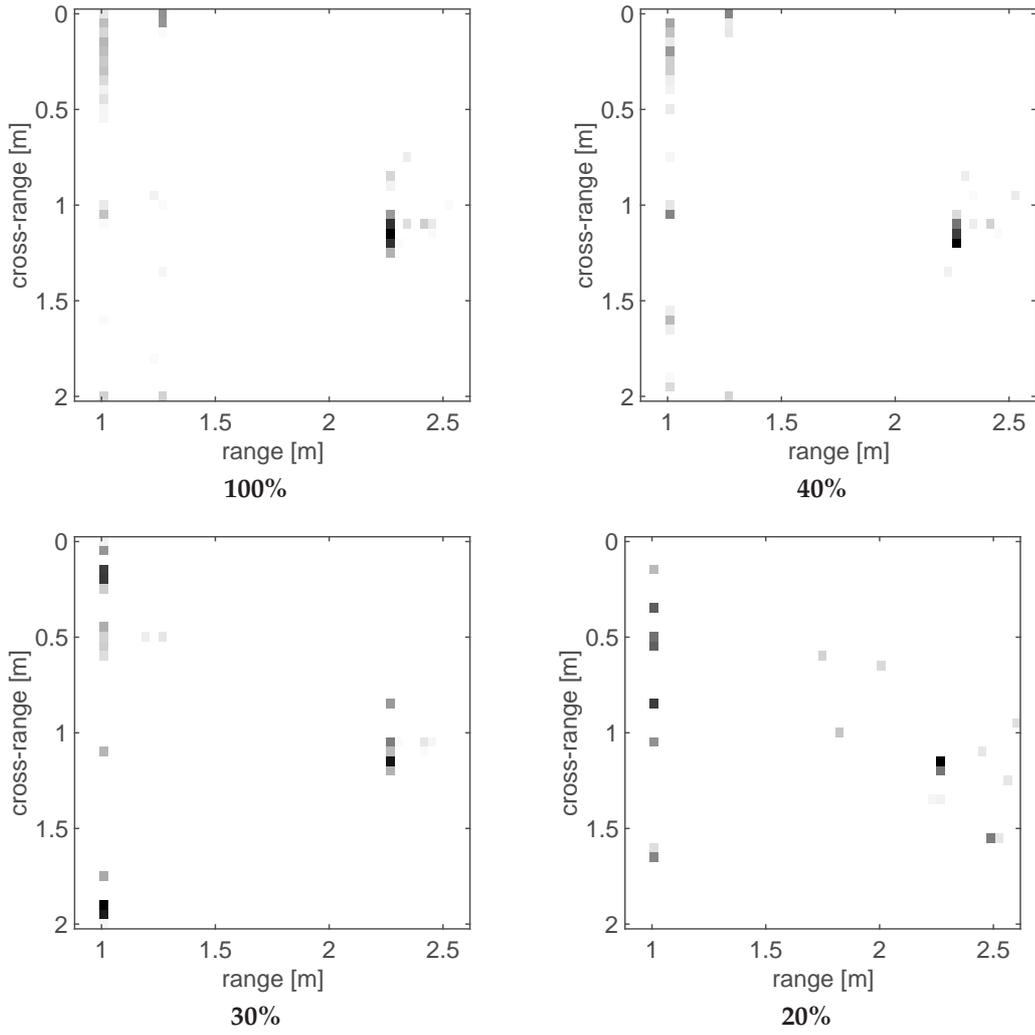


FIGURE 5.23: Reconstruction of the small scene with a single target using RPCA obtained using 100%, 40%, 30% and 20% of the original number of samples.

Target to Clutter Ratio				
SAR	BS	MA	SVD	RPCA
5.73 dB	32.93 dB	24.89 dB	26.69 dB	52.65 dB

5.4 Range profile reconstruction

In the previous sections of this chapter on TWI, we performed a complete SAR measurement and from the sub-Nyquist random sampled measurement we reconstructed the through the wall image using a CS algorithm. We will now apply a different approach to obtain the final image. First, we perform a measurement at each sensor position, subsequently we reconstruct the range compressed profile from the subsampled measurement at each position separately, using a CS reconstruction algorithm. Once we obtain all the range compressed profiles over all the sensor positions we reconstruct the fully compressed SAR image by applying a Range-Doppler cross-range compression (Figure 5.24).

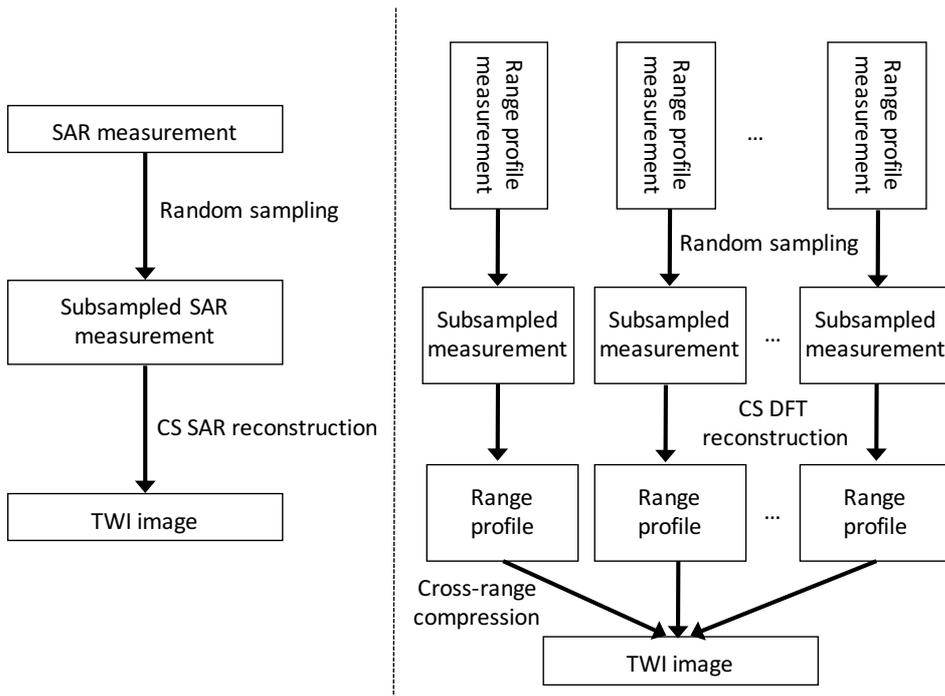


FIGURE 5.24: flowcharts of TWI CS approaches. On the left: CS SAR image reconstruction using a SAR measurement matrix. On the right: TWI obtained after a CS reconstruction of range profiles, using a DFT measurement matrix, and cross-range compression.

Our interest for studying this way of applying CS on TWI is motivated by following reflections:

1. The high similarity between the range profiles should be exploited during the SAR measurement phase. Once we have one range profile, this should be used to reconstruct the next one. With this approach we can mix the measurement phase and the

reconstruction phase which allows to use the reconstructed information during the remaining part of the SAR measurement.

2. The range profile reconstruction approach will allow to tailor the subsampling rate to the specific signal measured at that position and can thus vary along the scanning path.
3. Since the CS reconstruction is performed separately at each sensor position, we can select the same frequencies without compromising the final result. This can solve the problem described in subsection 5.3.3, when for example some frequencies are not allowed to use or are jammed.

5.4.1 Single range profile reconstruction

We will start with the range profile reconstruction obtained at a single sensor position. To do this, we concentrate on the range profile measured at sensor position 42, which corresponds to the sensor location at 2.10 m in the cross-range dimension and is right in front of the second human target. Figure 5.25 (a) shows the range compressed scene. The selected range position 42 is framed in red and the corresponding range profile is depicted in (b).

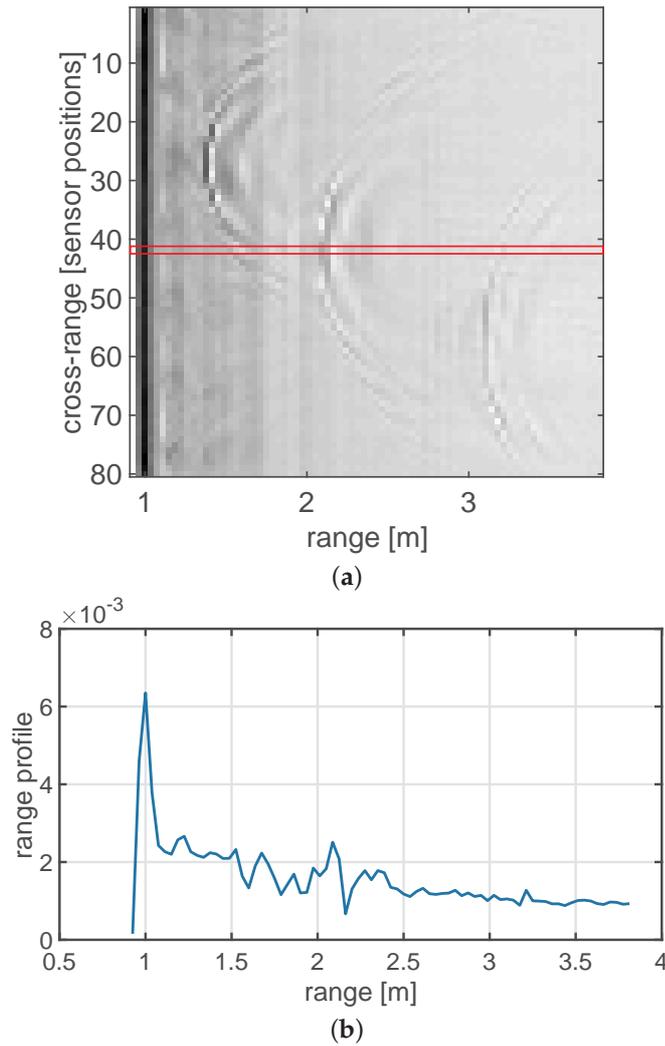


FIGURE 5.25: (a) Range compressed TWI with range range profile nr 42 framed with the red rectangle. (b) Plot of the range profile at position 42.

Figure 5.26 (a) shows the CS reconstruction of the selected range profile using 70%, 80%, 90% and 100% of the samples. The mean reconstruction error, calculated as the ℓ_2 -norm between the range profile obtained with 100% of the samples and the subsampled version is shown in Figure 5.26 (b). For each subsampling rate we performed 256 times a random sampling and reconstruction. We see an almost linear increase of the mean error for a decreasing subsampling rate.

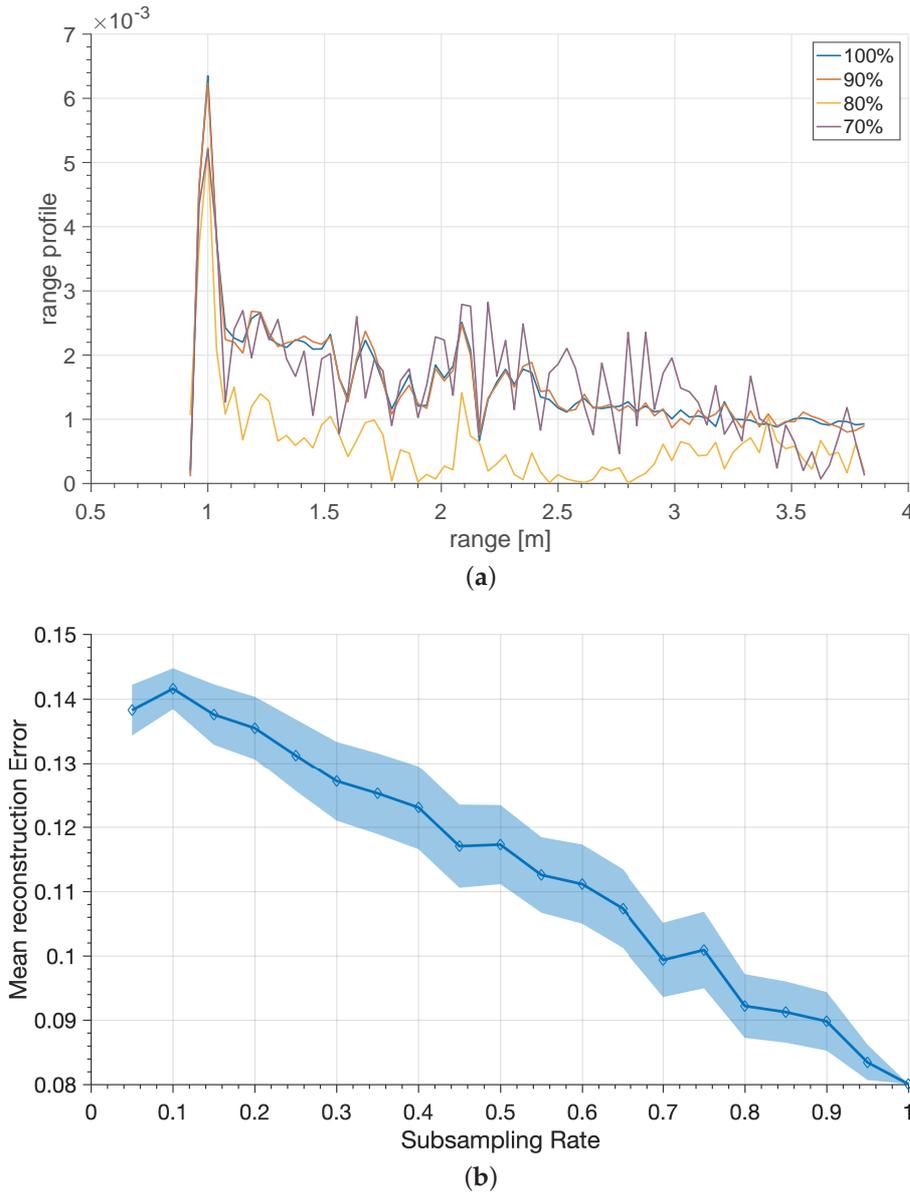


FIGURE 5.26: (a) CS reconstruction of range profile 42 using 100%, 90%, 80% and 70% of the original raw data. (b) Mean reconstruction error for increasing subsampling rates.

The same trend can be seen in Figure 5.27, where we depict the mean PSNR of the image obtained after CS SAR reconstruction (the blue curve) and the mean PSNR of the image obtained after range profile reconstruction and cross-range compression. Whereas the PSNR for the CS SAR reconstruction remains almost constant down to a subsampling rate of 75%, the PSNR of the reconstruction by the range profile CS approach immediately starts to drop.

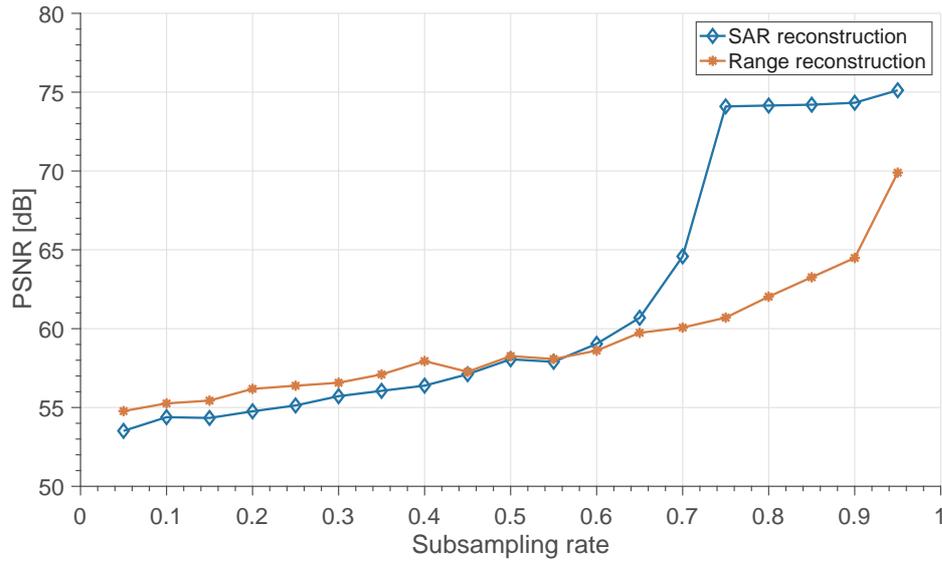


FIGURE 5.27: PSNR for the CS SAR image reconstruction (in blue) and the TWI obtained from the CS range profile reconstruction (in red) for increasing subsampling rates.

5.4.2 Side information

The mediocre results for the CS range profile reconstruction approach are explained by the fact that the range profiles are not sparse at all as can be seen in Figure 5.25. On the other hand we see a high similarity between the range profiles over the entire scanning distance of the sensor. We will now exploit this similarity by adding the previously measured range profiles as side information into the CS reconstruction. First we will do this by using a single range profile as side information and the reconstruction will be performed using an ℓ_1 - ℓ_1 minimization approach. Figure 5.28 (a) depicts the range profiles at sensor position 1 and position 42. We notice the similarity between the two profiles except for the zone where the reflections from the second target are received. As can be expected, the mean reconstruction error of the ℓ_1 - ℓ_1 approach is thus much lower (Figure 5.28) when SI (the range profile at sensor position 1) (red curve) is added, compared to the mean error for the reconstruction without SI (blue curve).

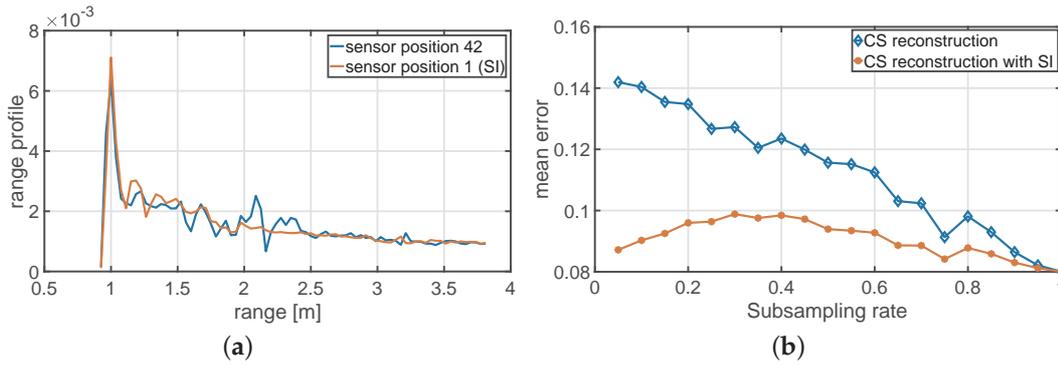


FIGURE 5.28: (a) Range profiles at sensor positions 1 (red curve) and 42 (blue curve). (b) Mean reconstruction error for the CS reconstruction of range profile 42 without SI (blue curve) and with SI (red curve).

5.4.3 Multiple weighted side information

The more relevant information that can be added during the CS reconstruction, the higher the robustness against subsampling. This means that not only a single range profile should be added as SI, but all the range profiles obtained at previous scanning positions should be exploited. Increasing the number of SIs also increases the risk of adding poor quality SI which can jeopardize the CS reconstruction. The RAMSIA algorithm faces this threat by granting inter and intra weights to the SIs. First: only those parts of the range profiles sharing a high degree of similarity with the range profile to reconstruct will be taken into account and second: the range profiles which are considered to be high quality SI will be granted more importance while the poor quality SI profiles will be neglected during the reconstruction.

Figure 5.29 (a) depicts the difference between the range profiles at scanning position 42 and position 1. This difference is much more sparse than the range profiles itself. In (b) we see the corresponding intra SI weights attributed automatically by the RAMSIA algorithm. We observe the inverse trend between the two curves. RAMSIA attributes high weights to those parts of range profile 1 where the difference with the range profile to reconstruct is minimal, whereas the elements with a poor similarity, for example the front wall reflection and the reflection of the second target (at a range distance of 2.25 m) will be suppressed.

Figure 5.30 illustrates the working of the inter-SI weights for the reconstruction of the range profile at the sensor position 78 (at 3.9 m from the starting position of the sensor). The graph in (a) depicts the difference (horizontal axis), calculated as the euclidean distance, between the range profile at sensor position 78 and all the previous range positions (vertical axis). The inter SI weights attributed by the RAMSIA algorithm to each of the range profiles used as SI is shown in (b). We observe again an inverse trend between the two graphs.

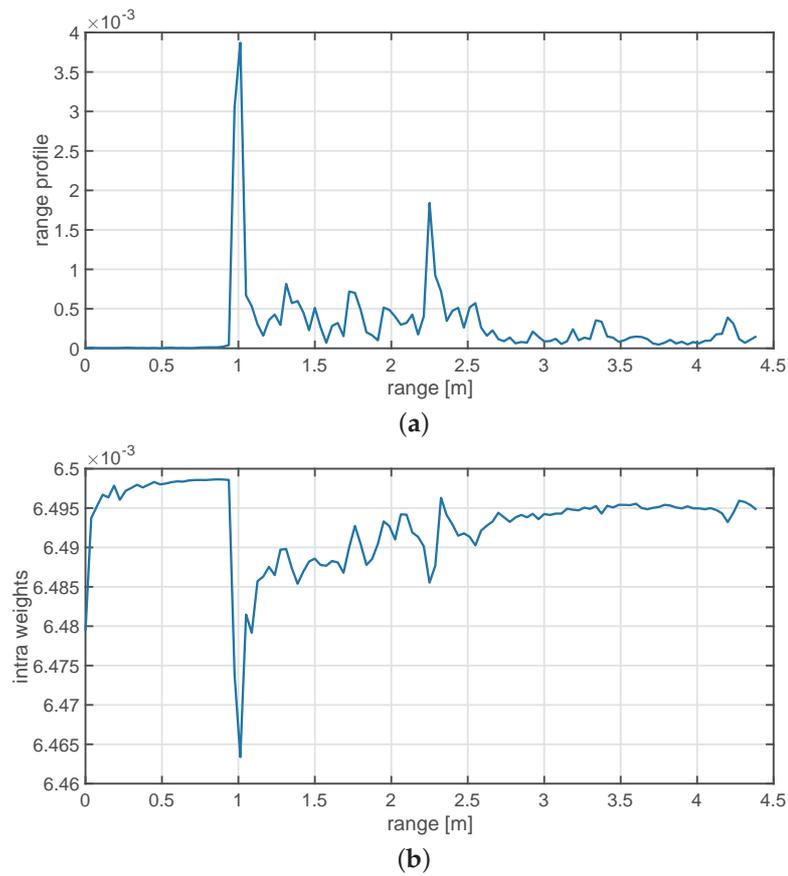


FIGURE 5.29: (a) Resulting range profile after subtracting the range profile at position 1 from the range profile at position 42. (b) Corresponding intra SI weights attributed by the RAMSIA algorithm.

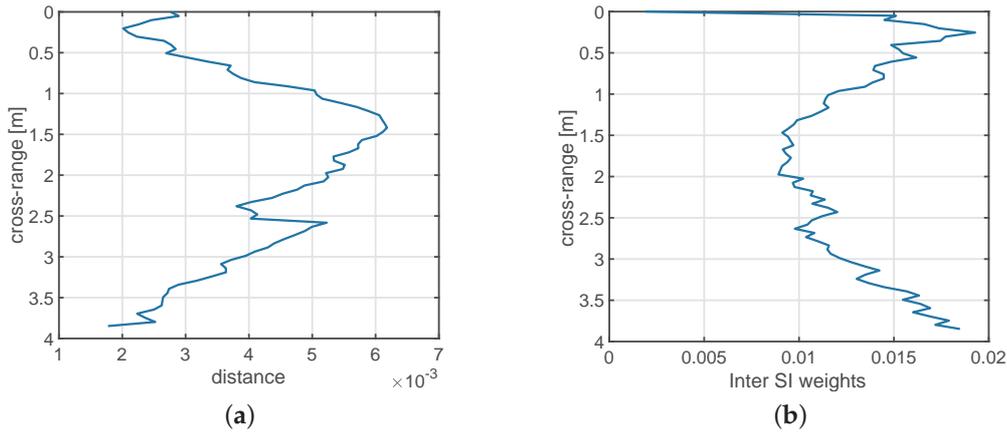


FIGURE 5.30: (a) Euclidean distance between the range profile at position 78 and all the previous range profiles. (b) attributed inter SI weights to the range profiles for the reconstruction of range profile 78.

The reconstruction of all the range profiles using RAMSIA allows to populate an ensemble of SIs which grows each time a new range profile is reconstructed while being almost immune against poor quality reconstructions. Figure 5.31 shows the resulting images for subsampling rates equal to 50%, 40%, 30% and 20%. We see that below 40% the targets start to vanish and at a subsampling rate of 20% the reconstruction quality is too low to detect any target.

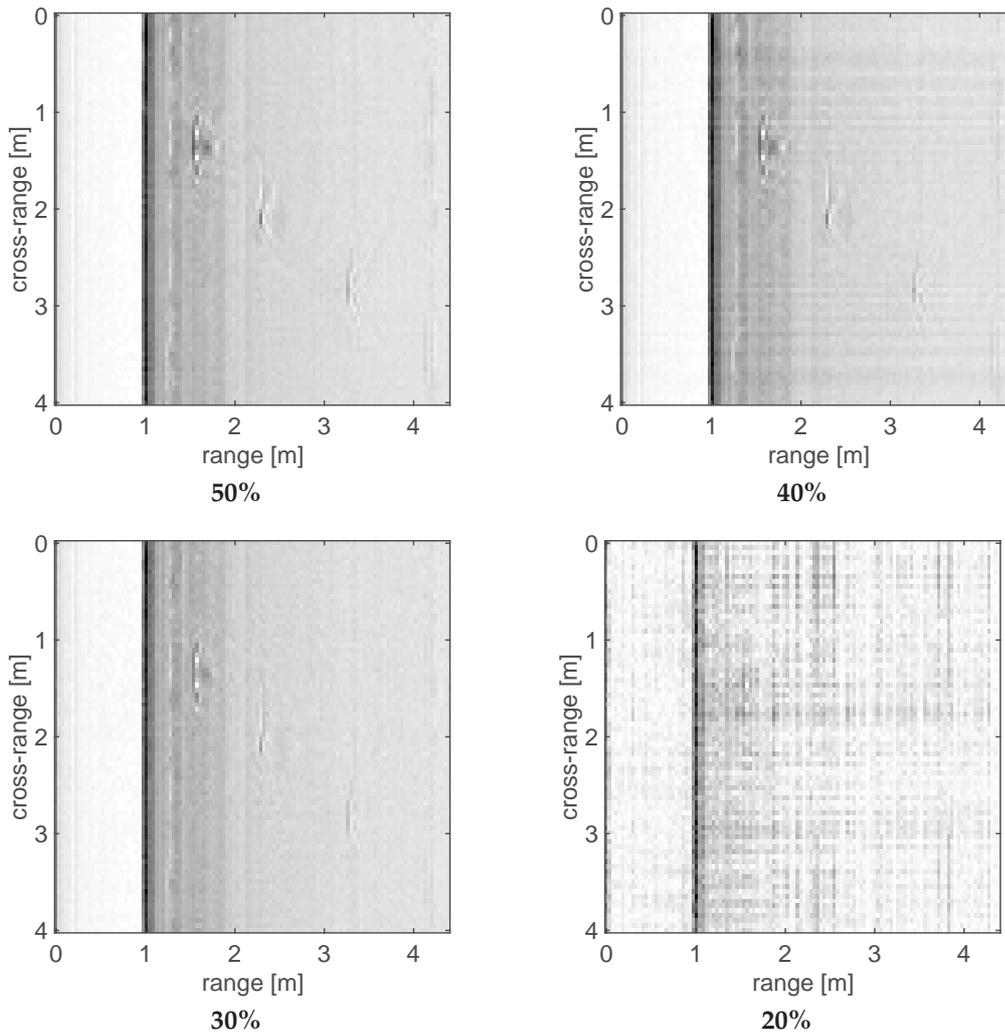


FIGURE 5.31: Resulting images after the reconstruction of the range profiles with RAMSIA for subsampling rates equal to 50%, 40%, 30% and 20%.

Figure 5.32 shows the reconstruction error (expressed in dB) for the reconstruction of each range profile along the scanning path (horizontal axis) as a function of the subsampling rate (vertical axis). In (a) we see the reconstruction error when the range profiles are reconstructed without SI and we observe that the error increases as soon as the number of samples start to decrease. In (b) we have the reconstruction errors obtained for the ℓ_1 - ℓ_1 CS reconstruction. Finally the best results, by far, are obtained in (c) when applying multiple weighted CS reconstruction.

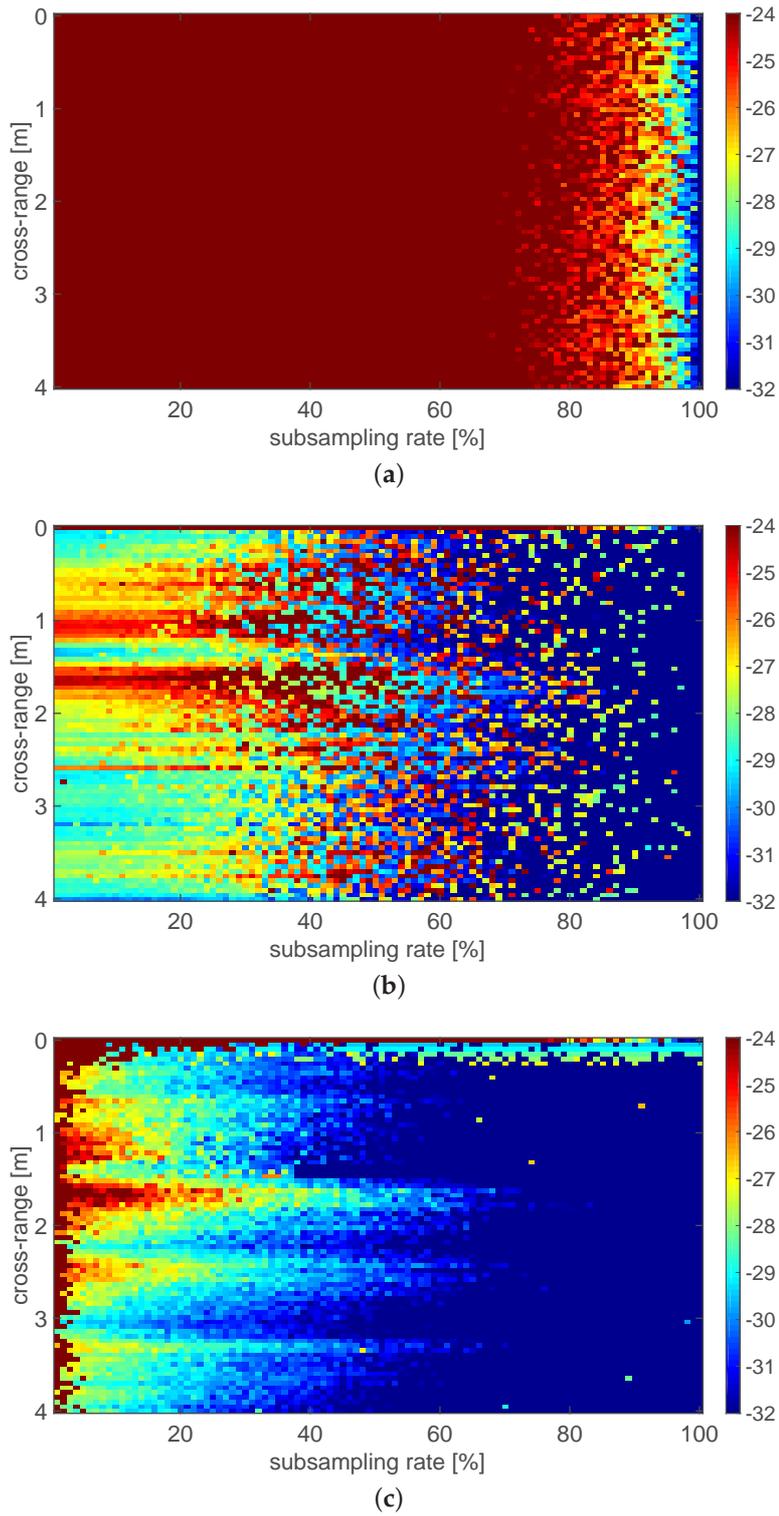


FIGURE 5.32: Reconstruction error [dB] of the range profiles over all sensor positions (vertical axis) for increasing subsampling rates (vertical axis). (a) without SI (b) $\ell_1\ell_1$ -reconstruction and (c) RAMSIA.

5.5 Sequential Compressed Sensing

Lots of research has been published on how to decrease the subsampling bound as much as possible. At the end of the day, this bound is a function of the sparsity of the signal to be reconstructed which, for most applications, is unknown prior to the measurement. In other words, we first need a measurement taken at the Nyquist rate in order to estimate the number of samples that are needed. A solution to this key problem is sequential compressed sensing and we will now evaluate the use of sequential compressed sensing combined with the CS reconstruction using multiple weighted side information for the reconstruction of TWI range profiles.

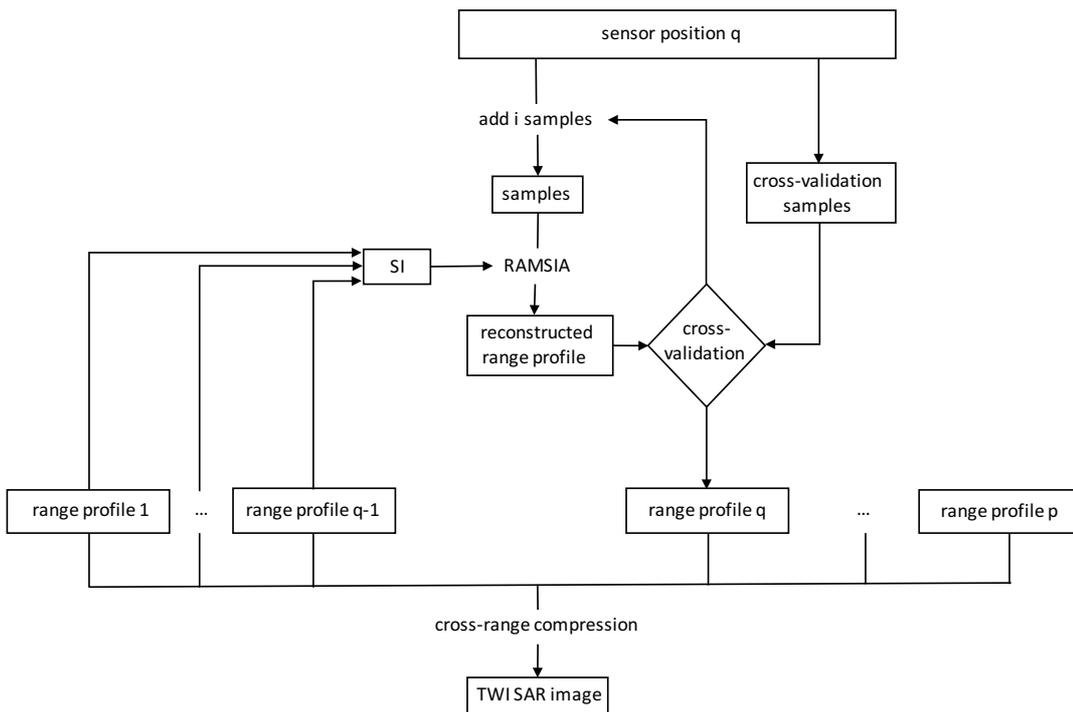


FIGURE 5.33: Flowchart explaining the approach combining multiple weighted side information and sequential compressed sensing.

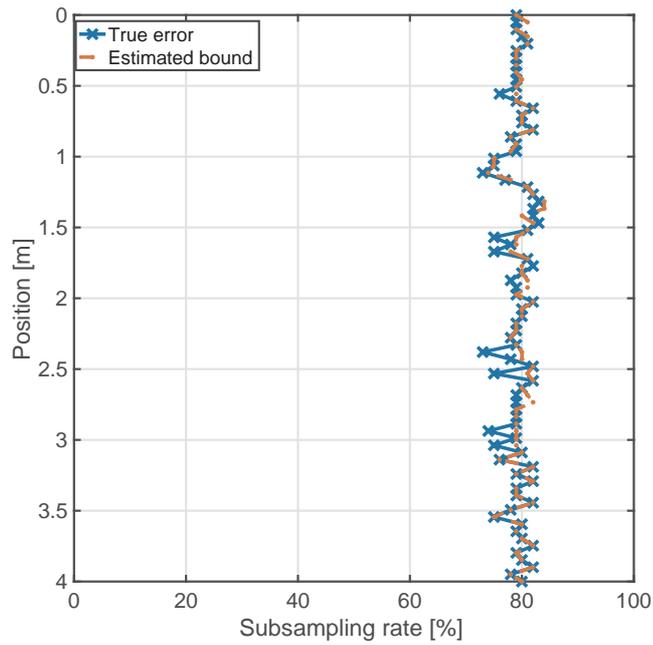
The flowchart in Figure 5.33 gives an overview of our approach at a single scanning position during the total SAR measurement. We notice that the proposed technique mixes the measurement phase and CS reconstruction phase. The CS reconstruction performed at a sensor position q will impact the measurements taken at the subsequent sensor positions by influencing the number of measurements that will be taken.

At scanning position q the approach starts by taking a small number of initial samples

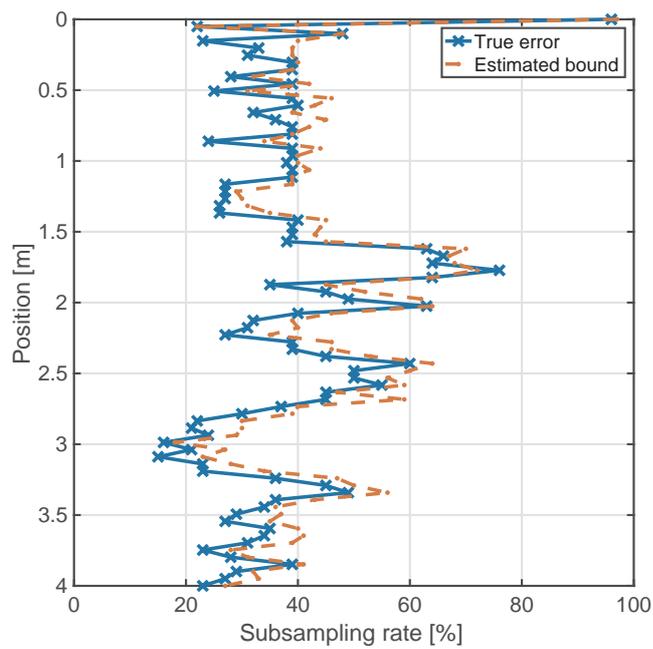
i and a small number t of extra samples which will serve as cross-validation measurements. Using the set of i samples, we reconstruct the range profile sensed from this position adding the range profiles obtained at all the previous scanning positions as weighted side information. The reconstructed range profile is then evaluated by using the cross-validation measurements. If the reconstruction does not pass the validation test, i new samples are added to the set of samples. This loop is repeated until the reconstructed range profile passes the validation test.

The validation test estimates an upper error bound for the reconstruction error. A good estimation of the error bound is thus necessary in order to avoid heavy oversampling or undersampling. Figure 5.33 depicts the true error (in blue) and the estimated error bounds (in red) for the reconstruction of all the range profiles of the measurement of the scene with the three human targets. In (a) we see the errors for the range profile reconstruction without the use of SI and in (b) we have the errors for the reconstruction using multiple weighted SI. In both cases the estimated bounds closely follow the actual error. These results were obtained by using 8 cross-validation samples. Figure 5.35 depicts the true errors and the corresponding estimated bounds for all range profiles reconstructed with the presented approach.

By combining the CS reconstruction of the range profiles using the other profiles as weighted SI with the sequential approach, we are able to build a sensor which decides autonomously on the number of samples that are taken in order to ensure a chosen reconstruction quality. The results depicted in Figure 5.36 show the reconstructed range profiles in (a) and the final obtained image after cross-range compression in (b). The only input to give is the desired reconstruction quality, which we chose arbitrarily equal to -30dB, the probability with which the quality will not drop below -30 dB, in our test equal to 95% and the number of cross-validation measurements, which we chose to be equal to 8. We have shown in chapter 3 (figure 3.18) that 8 cross-validation measurements are sufficient when choosing $\alpha = 0.05$. The algorithm automatically adds and weights the previous range profiles and adapts the number of samples at each position to assure the requested reconstruction quality. The number of samples taken at each sensor position are depicted in Figure 5.34. In order to obtain a maximum reconstruction error of -30 dB, 79.1% of the total number of samples are needed if no SI is used. By using weighted SI and allowing the number of samples to vary during the SAR measurement the number of samples can drop to 37.4%. When combining RAMSIA and using the sequential approach a reconstruction error of maximum -30 dB for each sensor position is guaranteed with a probability of 95% for a subsampling rate equal to 42.8%.



(a)



(b)

FIGURE 5.34: True error (in blue) and estimated error (in red) for in (a) using sequential CS without SI and in (b) using sequential CS with weighted SI.

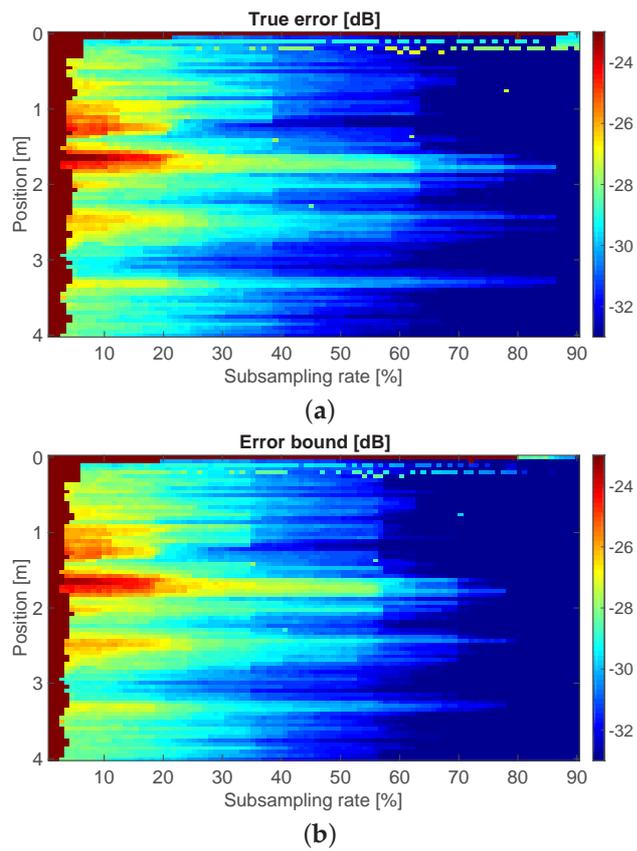


FIGURE 5.35: (a) True error [dB] and (b) the estimated error bounds [dB] for the reconstruction of all the range profiles along the scanning path for increasing subsampling rates.

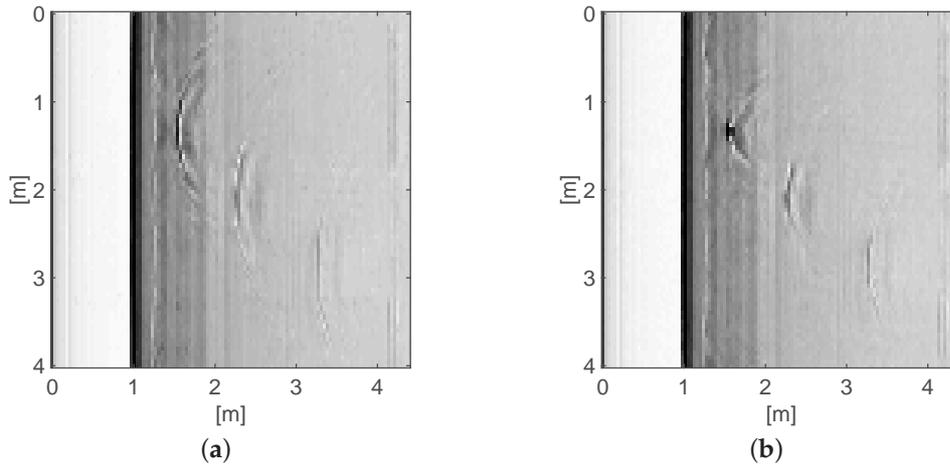


FIGURE 5.36: (a) the reconstructed range profiles and (b) the reconstructed TWI SAR image by combining RAMSIA and sequential CS.

5.6 Conclusions

This chapter started with a literature study and compared the state of the art research in TWI with the commercial-of-the-shelf TWI sensors. A large gap exists between those sensors. Research prototypes deliver good quality images of the scenes behind the wall but require high bandwidths, large antenna arrays or scanning paths, whereas the commercial systems are small and fast but deliver only low quality images. Compressed sensing can help to close the gap, by reducing the number of samples, the number of scanning positions and the measurement time.

A TWI SAR prototype was built in order to perform experiments and acquire real TWI radar data. The system was used to image scenes behind a single layered wall populated with multiple static human targets. The resulting SAR images are contaminated with wall clutter making the detection of the targets difficult. Moreover, the obtained images are not sparse and thus sensitive to subsampling when applying a CS reconstruction approach. Two possible solutions were tested. (1) We utilized a discrete wavelet transform and a DCT as sparsifying bases and (2) we tested four wall mitigation techniques (coherent background subtraction, moving average filter, singular value decomposition and robust principal component analysis). We obtained good results by applying the DCT as a sparsifying basis. All of the four tested wall mitigation approaches increase the TCR but only the background subtraction and RPCA approach are robust against subsampling.

We explored an alternative approach for applying CS by reconstructing the range profiles instead of reconstructing directly the image. The strategy allows to add the previous reconstructed range profiles as weighted side information. We further tested a sequential CS approach. By combining multiple weighted CS and sequential CS we were able to conceive and test a sensor which automatically determines the number of samples to

take while performing the measurement and at the same time uses that part of the SAR measurement which has already been executed to lower the number of measurements.

Chapter 6

Conclusions

In the final chapter of this dissertation, we summarize the most important results of each chapter individually. We further draw the general conclusions on our study on applying Compressed Sensing on in-depth microwave imaging. We will end this chapter with a series of ideas for future work.

6.1 Conclusions of the different chapters

Chapter 2: Compressed Sensing

Chapter 2 started with a concise overview of the CS theory. The basic idea of CS is to lower the subsampling bound, dictated by the Shannon-Nyquist theorem, by exploiting the prior knowledge that the sensed signal is sparse in some domain. The given theoretical results on CS showed that CS can decrease the sampling bound drastically if the measurement modality satisfies some property. Unfortunately, these properties are difficult to verify or lead to suboptimal bounds. On the other hand, effective algorithms for reconstructing sparse signals are well known. In this chapter, we explored two families of algorithms: Basis Pursuit algorithms and (2) Matching Pursuit algorithms. We compared several algorithms (which are used in the later chapters of this dissertation) from both families in terms of running time, uniformity and robustness against subsampling. The COSAMP algorithm revealed to be fast compared to the Basis Pursuit algorithms and uniform compared to OMP. The Basis Pursuit algorithms, on the other hand, are robust against subsampling and are also uniform.

Chapter 3: Compressed Sensing applied on SFCW and SAR measurements

In chapter 3, we concentrated on using CS on Stepped-Frequency Continuous Wave (SFCW) measurements and on Synthetic Aperture Radar (SAR) measurements. Through a series of synthetic measurements, we demonstrated that CS allows to significantly decrease the number of samples while preserving the quality of the reconstructed signal. The new lower bound is a function of the a priori unknown sparsity of the sensed signal. The lower bound on the number of samples is thus unknown prior to the measurement.

We demonstrated that we can estimate a sharp bound on the reconstruction error by using only a small number of cross-validation samples. This technique can be used as a stopping criterion in a configuration where we sequentially add samples without having access to the ground truth.

In this chapter, we further introduced different techniques which can be combined with CS to further lower the CS sampling bound: the Singular Value Decomposition and the Principal Component Analysis.

We ended chapter 3 with a study on adding Side Information (SI) to the CS minimization problem. We first compared different approaches to add homogeneous SI. The RAMSIA algorithm, which adds weights to the SI, revealed to outperform the other approaches, such as the coherent background subtraction. We further explored a way to add homogeneous SI. The proposed approach consists of two steps: (1) a coupled dictionary learning step and (2) a CS reconstruction step using the heterogeneous SI. We showed that the technique allows to perform change detection from heavily subsampled measurements by exploiting the information coming from a different sensor.

The discussed techniques in this chapter are further used and evaluated on the real data in chapters 4 and 5 of this dissertation.

Chapter 4: Non-Destructive testing of polymer 3-D printed objects

In chapter 4, we evaluated the applicability of CS and CS with SI on real SAR data. The first application is the Non-Destructive Evaluation (NDE) of 3-D printed parts using a mm-wave sensor. The chapter starts with a general introduction on 3-D printing and an overview of the pros and cons of different existing NDE technologies. To obtain the real NDE data which we used throughout chapter 4, we used a mm-wave sensor based on a MNVA. This technique has the advantage that it uses non-ionising radiation and that the measurements are performed in a stand-off setup. In order to obtain an image with an acceptable resolution, a large bandwidth is needed, combined with a large (synthetic) antenna. The measurements process is thus time consuming and produces a vast amount of raw data. In this chapter, we showed that CS can help to reduce the large amount of samples through a series of experiments on a 3-D printed object with synthetically created defects. We further compared three different techniques to add homogeneous SI to the CS reconstruction: (1) Background subtraction, (2) $\ell_1\ell_1$ -minimization and (3) CS with weighted SI. Over all the tests performed on the real NDT data, the best results were obtained by applying CS with multiple weighted SI. The introduced defects are clearly visible on the image obtained with RAMSIA using only 30% of the samples.

Chapter 5: Through-the-Wall Imaging

We started chapter 5 with a literature study of the research TWI prototypes and the commercially available TWI sensors. The comparison revealed that, on the one hand, the research prototypes deliver high resolution images but need large antenna arrays or synthetically created large antennas and emits wideband signals. The operational sensors, on the other hand, are lightweight and easy to deploy but do not deliver high quality images. In this chapter, we explored if it is possible to combine the best of both worlds, by using Compressed Sensing techniques. To do so, we built a TWI sensor prototype and performed a series of TWI measurements.

The first tests we performed using CS did not deliver hopeful results: only a decrease of

30% of the original number of samples was possible. This is due to the fact that the image is corrupted by the effects of the wall and is not sparse. A first solution to this problem is the use of a sparsifying transform. By reconstructing the Discrete Cosine Transform of the image, the subsampling rate can be lowered down to 25%. We also evaluated different existing techniques for mitigating the front wall clutter in combination with CS: (1) Coherent background subtraction, (2) Moving average filter, (3) Singular Value Decomposition (SVD) and (4) Robust Principal Component Analysis (RPCA). The different techniques decrease the Target-to-Clutter Ratio (TCR) but only the background subtraction and RPCA approaches are robust against subsampling.

We ended the chapter by testing a novel approach. Instead of reconstructing the SAR TWI image at the end of the measurement phase, we reconstructed the current range profile before moving the sensor to the next scanning position. The already measured range profiles are added as weighted SI to the reconstruction of the range profile. We combined this approach with a Sequential CS approach. This resulted in a TWI sensor which adapts autonomously the number of samples along the scanning path without having access to the ground truth and assures to deliver an image with a chosen reconstruction quality. This implies that the operator no longer has to guess the sparsity of the scene in order to decide on the number of samples. This new approach was tested on the real data and allows a decrease of the number of samples down to 42.8%.

6.2 Final conclusion

At the end of this work we can conclude that CS is a good candidate for lowering the number of samples for sensors performing mm-wave SAR measurements. This technique can be combined with other signal processing techniques to further lower the sampling bound. We have extensively shown that adding weighted Side Information to the CS reconstruction can optimise the working of the sensor. The techniques were successfully tested using real measurements in the domains of Non-Destructive Evaluation of 3-D printed objects and Through-the-Wall Imaging. I strongly believe that these results bring the mm-wave sensors closer to an operational level in these domains.

6.3 Future work

We will now formulate a series of hints and ideas for future work.

- Future work on mm-wave NDT of 3-D printed objects:
 1. In chapter 4, we performed a series of experiments on a simple test object with synthetically created defects. The NDT results should be validated with experiments on objects with a more complex geometry and containing real defects. These results should then be compared with the results obtained from NDT measurements on the same objects with different NDT techniques.
 2. The most important drawback of mm-wave imaging used as a NDT technique is the mediocre resolution of the images when compared to X-ray imaging. In

order to improve the resolution, the center frequency should increase and the total emitted bandwidth should be enlarged.

3. In section 3.6.2, we analyzed a technique to add heterogeneous SI to the CS reconstruction of a mm-wave SAR measurement. The results obtained on synthetic data were promising. Because of the high number of learning measurements required to obtain the coupled dictionaries, it was not feasible to perform enough real tests with two different sensors within the framework of this project. Testing this approach on real NDE data by combining two different measurement modalities (for example: X-ray and mm-wave measurements) is definitely a promising and interesting challenge.
- Future work on SAR Through-the-Wall Imaging:
 1. In section 5.3.4, two different sparsifying transforms were tested: the DCT transform and the Haar wavelet. This study can be enlarged by evaluating other transforms which can lead to even sparser representations of the TWI image. Alternatively, a dictionary can be learned from the TWI measurements. This dictionary can even further be optimized after each measurement.
 2. In chapter 5, the already reconstructed range profiles are added as weighted SI to the CS reconstruction of the current range profile. Alternatively, the Compressive Online Robust Principal Component Analysis (CORPCA) algorithm [Van Luong et al.2017] could be used. This approach performs an online separation of the sparse foreground and the low-rank background using the results obtained at the previous scanning positions. This technique should improve the Target-to-Clutter Ratio (TCR) and at the same time being robust against sub-sampling. Unfortunately, we were not able (yet) to combine this technique with sequential CS.
 - Future work on CS for other applications:

CS applied to SFCW or SAR measurements using heterogeneous SI could be a solution in other applications where the transmission of the raw data is costly or where the data of different sensors could be combined. For example:

 1. For the transmission of SAR data which are acquired on board of an airborne platform or a satellite. In most cases, reliable visual images of the scene are available (e.g.: Google Earth) and can be used as heterogeneous SI. Exploiting the heterogeneous SI, could drastically reduce the data volumes that need to be transmitted since most of the transmitted information is already known.
 2. Another possible application of coupled dictionary learning can be found in the automotive industry. The data obtained by one sensor can serve to reduce the acquisition time or to fill in the gapped data acquired by a different sensor.

Bibliography

- [SPG] <http://www.cs.ubc.ca/mpf/spgl1/>.
- [Abdeladi2015] Abdeladi, A. (2015). Multipath exploitation in through the wall radar imaging using sparsity driven detection. Master's thesis.
- [Acharya and Tsai2005] Acharya, T. and Tsai, P.-S. (2005). *JPEG2000 standard for image compression: concepts, algorithms and VLSI architectures*. John Wiley & Sons.
- [Aharon et al.2006] Aharon, M., Elad, M., Bruckstein, A., et al. (2006). K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing*, 54(11):4311.
- [Ahmad2008] Ahmad, F. (2008). Multi-location wideband through-the-wall beamforming. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 5193–5196. IEEE.
- [Ahmad and Amin2012] Ahmad, F. and Amin, M. G. (2012). Partially sparse reconstruction of behind-the-wall scenes. In *Compressive Sensing*, volume 8365, page 83650W. International Society for Optics and Photonics.
- [Ahmad and Amin2013] Ahmad, F. and Amin, M. G. (2013). Through-the-wall human motion indication using sparsity-driven change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 51(2):881–890.
- [Ahmad et al.2013] Ahmad, F., Qian, J., and Amin, M. G. (2013). Wall mitigation using discrete prolate spheroidal sequences for sparse indoor image reconstruction. In *Signal Processing Conference (EUSIPCO), 2013 Proceedings of the 21st European*, pages 1–5. IEEE.
- [Ahmad et al.2015] Ahmad, F., Qian, J., and Amin, M. G. (2015). Wall clutter mitigation using discrete prolate spheroidal sequences for sparse reconstruction of indoor stationary scenes. *IEEE Transactions on Geoscience and Remote Sensing*, 53(3):1549–1557.
- [Ahmed et al.2012] Ahmed, S. S., Schiessl, A., Gumbmann, F., Tiebout, M., Methfessel, S., and Schmidt, L.-P. (2012). Advanced microwave imaging. *IEEE microwave magazine*, 13(6):26–43.
- [Amin2017] Amin, M. G. (2017). *Through-the-wall radar imaging*. CRC press.
- [Amin and Ahmad2014] Amin, M. G. and Ahmad, F. (2014). Through-the-wall radar imaging: theory and applications. In *Academic Press Library in Signal Processing*, volume 2, pages 857–909. Elsevier.
- [Bamler1992] Bamler, R. (1992). A comparison of range-Doppler and wavenumber domain SAR focusing algorithms. *IEEE Transactions on Geoscience and Remote Sensing*, 30(4):706–713.
- [Bandeira et al.2013] Bandeira, A. S., Dobriban, E., Mixon, D. G., and Sawin, W. F. (2013). Certifying the restricted isometry property is hard. *IEEE transactions on information theory*, 59(6):3448–3450.

- [Baraniuk et al.2008] Baraniuk, R., Davenport, M., DeVore, R., and Wakin, M. (2008). A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263.
- [Barrie2004] Barrie, G. B. (2004). Through-wall synthetic aperture radar (TWSAR) 3D imaging: algorithm design. Technical report, DEFENCE RESEARCH AND DEVELOPMENT CANADA OTTAWA (ONTARIO).
- [Beck and Teboulle2009] Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202.
- [Bertsekas2014] Bertsekas, D. P. (2014). *Constrained optimization and Lagrange multiplier methods*. Academic press.
- [Blitz2012] Blitz, J. (2012). *Electrical and magnetic methods of non-destructive testing*, volume 3. Springer Science & Business Media.
- [Borek2005] Borek, S. E. (2005). An overview of through the wall surveillance for homeland security. In *Applied Imagery and Pattern Recognition Workshop, 2005. Proceedings. 34th*, pages 6–pp. IEEE.
- [Bouzerdoun and Tivive2015] Bouzerdoun, A. and Tivive, F. H. C. (2015). Wall clutter mitigation using HOSVD in through-the-wall radar imaging with compressed sensing. In *Digital Signal Processing (DSP), 2015 IEEE International Conference on*, pages 85–89. IEEE.
- [Boyd and Vandenberghe2004] Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- [Campbell et al.2011] Campbell, T., Williams, C., Ivanova, O., and Garrett, B. (2011). Could 3D printing change the world. *Technologies, Potential, and Implications of Additive Manufacturing*, Atlantic Council, Washington, DC.
- [Candes and Romberg2005] Candes, E. and Romberg, J. (2005). l1-magic: Recovery of sparse signals via convex programming. URL: www.acm.caltech.edu/l1magic/downloads/l1magic.pdf, 4:14.
- [Candes2008] Candes, E. J. (2008). The restricted isometry property and its implications for compressed sensing. *Comptes rendus mathématique*, 346(9-10):589–592.
- [Candès et al.2011] Candès, E. J., Li, X., Ma, Y., and Wright, J. (2011). Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11.
- [Candès et al.2006] Candès, E. J., Romberg, J., and Tao, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509.
- [Candes et al.2006] Candes, E. J., Romberg, J. K., and Tao, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223.
- [Candes and Tao2005] Candes, E. J. and Tao, T. (2005). Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215.
- [Cao et al.2016] Cao, X., Yang, L., and Guo, X. (2016). Total variation regularized RPCA for irregularly moving object detection under dynamic background. *IEEE transactions on cybernetics*, 46(4):1014–1027.

- [Cevher et al.2008] Cevher, V., Sankaranarayanan, A., Duarte, M. F., Reddy, D., Baraniuk, R. G., and Chellappa, R. (2008). Compressive sensing for background subtraction. In *European Conference on Computer Vision*, pages 155–168. Springer.
- [Chandrasekaran et al.2011] Chandrasekaran, V., Sanghavi, S., Parrilo, P. A., and Willsky, A. S. (2011). Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596.
- [Cristofani et al.2018] Cristofani, E., Becquaert, M., Lambot, S., Vandewal, M., Stiens, J., and Deligiannis, N. (2018). Random subsampling and data preconditioning for ground penetrating radars. *IEEE Access*.
- [Dasgupta and Gupta2003] Dasgupta, S. and Gupta, A. (2003). An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65.
- [Davenport et al.2011] Davenport, M. A., Duarte, M. F., Eldar, Y. C., and Kutyniok, G. (2011). Introduction to compressed sensing. *preprint*, 93(1):2.
- [de Wit et al.2006] de Wit, J. J., Meta, A., and Hooeboom, P. (2006). Modified range-Doppler processing for FM-CW synthetic aperture radar. *IEEE Geoscience and Remote Sensing Letters*, 3(1):83–87.
- [Debes et al.2010] Debes, C., Riedler, J., Zoubir, A. M., and Amin, M. G. (2010). Adaptive target detection with application to through-the-wall radar imaging. *IEEE Transactions on Signal Processing*, 58(11):5572–5583.
- [Dehmollaian and Sarabandi2008] Dehmollaian, M. and Sarabandi, K. (2008). Refocusing through building walls using synthetic aperture radar. *IEEE Transactions on Geoscience and Remote Sensing*, 46(6):1589–1599.
- [Deligiannis et al.2017] Deligiannis, N., Mota, J. F., Cornelis, B., Rodrigues, M. R., and Daubechies, I. (2017). Multi-modal dictionary learning for image separation with application in art investigation. *IEEE Transactions on Image Processing*, 26(2):751–764.
- [Ding et al.2017] Ding, X., Li, H., Zhu, J., Wang, G., Cao, H., Zhang, Q., and Ma, H. (2017). Application of infrared thermography for laser metal-wire additive manufacturing in vacuum. *Infrared Physics & Technology*, 81:166–169.
- [Donoho and Tanner2009] Donoho, D. and Tanner, J. (2009). Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 367(1906):4273–4293.
- [Donoho2006] Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306.
- [Donoho et al.2012] Donoho, D. L., Tsaig, Y., Drori, I., and Starck, J.-L. (2012). Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 58(2):1094–1121.
- [du Plessis et al.2016] du Plessis, A., le Roux, S. G., Booysen, G., and Els, J. (2016). Quality control of a laser additive manufactured medical implant by X-ray tomography. *3D Printing and Additive Manufacturing*, 3(3):175–182.
- [Duarte and Eldar2011] Duarte, M. F. and Eldar, Y. C. (2011). Structured compressed sensing: From theory to applications. *IEEE Transactions on Signal Processing*, 59(9):4053–4085.

- [Eldar et al.2010] Eldar, Y., Kuppinger, P., and Bolcskei, H. (2010). Compressed sensing of blocksparse signals: uncertainty relations and efficient recovery. submitted to *IEEE Transactions on Signal Processing*.
- [Ertin and Moses2009] Ertin, E. and Moses, R. L. (2009). Through-the-wall SAR attributed scattering center feature estimation. *IEEE Transactions on Geoscience and Remote Sensing*, 47(5):1338–1348.
- [Everton et al.2015] Everton, S., Dickens, P., Tuck, C., and Dutton, B. (2015). Evaluation of laser ultrasonic testing for inspection of metal additive manufacturing. In *Laser 3d Manufacturing II*, volume 9353, page 935316. International Society for Optics and Photonics.
- [Everton et al.2016] Everton, S. K., Hirsch, M., Stravroulakis, P., Leach, R. K., and Clare, A. T. (2016). Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing. *Materials & Design*, 95:431–445.
- [Farwell et al.2008] Farwell, M., Ross, J., Luttrell, R., Cohen, D., Chin, W., and Dogaru, T. (2008). Sense through the wall system development and design considerations. *Journal of the Franklin Institute*, 345(6):570–591.
- [Forrest et al.1996] Forrest, J., Dalnoki-Veress, K., Stevens, J., and Dutcher, J. (1996). Effect of free surfaces on the glass transition temperature of thin polymer films. *Physical Review Letters*, 77(10):2002.
- [Foucart and Rauhut2013] Foucart, S. and Rauhut, H. (2013). *A mathematical introduction to compressive sensing*, volume 1. Birkhäuser Basel.
- [Frazier1997] Frazier, L. M. (1997). Radar surveillance through solid materials. In *Command, Control, Communications, and Intelligence Systems for Law Enforcement*, volume 2938, pages 139–147. International Society for Optics and Photonics.
- [Gaikwad and Shevada2013] Gaikwad, A. N. and Shevada, L. K. (2013). Study of compressive sensing on through wall imaging. *International Journal of Computer Applications*, 72(21).
- [Gamper et al.2008] Gamper, U., Boesiger, P., and Kozerke, S. (2008). Compressed sensing in dynamic MRI. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 59(2):365–373.
- [Gosselin et al.2016] Gosselin, C., Duballet, R., Roux, P., Gaudillière, N., Dirrenberger, J., and Morel, P. (2016). Large-scale 3D printing of ultra-high performance concrete—a new processing route for architects and builders. *Materials & Design*, 100:102–109.
- [Grant et al.2008] Grant, M., Boyd, S., and Ye, Y. (2008). CVX: Matlab software for disciplined convex programming.
- [Gunatilaka and Baertlein2000] Gunatilaka, A. H. and Baertlein, B. A. (2000). Subspace decomposition technique to improve GPR imaging of antipersonnel mines. In *Detection and Remediation Technologies for Mines and Minelike Targets V*, volume 4038, pages 1008–1019. International Society for Optics and Photonics.
- [Guo et al.2018] Guo, S., Yang, X., Cui, G., Song, Y., and Kong, L. (2018). Multipath ghost suppression for through-the-wall imaging radar via array rotating. *IEEE Geoscience and Remote Sensing Letters*, 15(6):868–872.
- [Haldar et al.2011] Haldar, J. P., Hernando, D., and Liang, Z.-P. (2011). Compressed-sensing MRI with random encoding. *IEEE transactions on Medical Imaging*, 30(4):893–903.

- [Hassen and Kirka2018] Hassen, A. A. and Kirka, M. M. (2018). Additive manufacturing: The rise of a technology and the need for quality control and inspection techniques. *Materials Evaluation*, 76(4):438–453.
- [Hassen et al.2016] Hassen, A. A., Lindahl, J., Chen, X., Post, B., Love, L., and Kunc, V. (2016). Additive manufacturing of composite tooling using high temperature thermoplastic materials. In *SAMPE Conference Proceedings, Long Beach, CA, May*, pages 23–26.
- [Hausman and Horne2014] Hausman, K. K. and Horne, R. (2014). *3D printing for dummies*. John Wiley & Sons.
- [Helander et al.2017] Helander, J., Ericsson, A., Gustafsson, M., Martin, T., Sjöberg, D., and Larsson, C. (2017). Compressive sensing techniques for mm-wave nondestructive testing of composite panels. *IEEE Transactions on Antennas and Propagation*, 65(10):5523–5531.
- [Hong et al.2011] Hong, M., Yu, Y., Wang, H., Liu, F., and Crozier, S. (2011). Compressed sensing MRI with singular value decomposition-based sparsity basis. *Physics in Medicine & Biology*, 56(19):6311.
- [Hull1986] Hull, C. W. (1986). Apparatus for production of three-dimensional objects by stereolithography. US Patent 4,575,330.
- [Jerr1977] Jerri, A. J. (1977). The Shannon sampling theorem—its various extensions and applications: A tutorial review. *Proceedings of the IEEE*, 65(11):1565–1596.
- [Jet Propulsion Laboratory2015] Jet Propulsion Laboratory, C. I. o. T. (2015). DHS and NASA technology helps save four in Nepal earthquake disaster. Available: <http://www.jpl.nasa.gov/news/news.php?feature=4574>.
- [Katz1989] Katz, N. M. (1989). An estimate for character sums. *Journal of the American Mathematical Society*, 2(2):197–200.
- [Khamlichi et al.2012] Khamlichi, J., Chappe, O., Lescalier, B., Gaugue, A., Menard, M., et al. (2012). Using CFAR algorithm to further improve a combined through-wall imaging method. In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, pages 2521–2525. IEEE.
- [Kharkovsky et al.2006] Kharkovsky, S., Case, J. T., Abou-Khousa, M. A., Zoughi, R., and Hepburn, F. L. (2006). Millimeter-wave detection of localized anomalies in the space shuttle external fuel tank insulating foam. *IEEE Transactions on Instrumentation and Measurement*, 55(4):1250–1257.
- [Kharkovsky and Zoughi2007] Kharkovsky, S. and Zoughi, R. (2007). Microwave and millimeter wave nondestructive testing and evaluation-overview and recent advances. *IEEE Instrumentation & Measurement Magazine*, 10(2):26–38.
- [Koester et al.2018] Koester, L. W., Taheri, H., Bigelow, T. A., Collins, P. C., and Bond, L. J. (2018). Nondestructive testing for metal parts fabricated using powder-based additive manufacturing. *Materials Evaluation*, 76(4):514–524.
- [Kotelnikov1933] Kotelnikov, V. A. (1933). On the transmission capacity of the ‘ether’ and of cables in electrical communications. In *Proceedings of the first All-Union Conference on the technological reconstruction of the communications sector and the development of low-current engineering*. Moscow. Citeseer.
- [Lagunas et al.2012a] Lagunas, E., Amin, M. G., Ahmad, F., and Najar, M. (2012a). Compressive sensing for through wall radar imaging of stationary scenes using arbitrary

- data measurements. In *Information Science, Signal Processing and their Applications (ISSPA), 2012 11th International Conference on*, pages 1347–1352. IEEE.
- [Lagunas et al.2012b] Lagunas, E., Amin, M. G., Ahmad, F., and Nájjar, M. (2012b). Wall mitigation techniques for indoor sensing within the compressive sensing framework. In *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2012 IEEE 7th*, pages 213–216. IEEE.
- [Le Gall1991] Le Gall, D. (1991). MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):46–58.
- [Leigsnering et al.2014] Leigsnering, M., Ahmad, F., Amin, M., and Zoubir, A. (2014). Multipath exploitation in through-the-wall radar imaging using sparse reconstruction. *IEEE Transactions on Aerospace and Electronic Systems*, 50(2):920–939.
- [Leigsnering et al.2011] Leigsnering, M., Debes, C., and Zoubir, A. M. (2011). Compressive sensing in through-the-wall radar imaging. In *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, pages 4008–4011. IEEE.
- [Lewis and Knowles1992] Lewis, A. S. and Knowles, G. (1992). Image compression using the 2-D wavelet transform. *IEEE Transactions on image Processing*, 1(2):244–250.
- [Liebgott et al.2013] Liebgott, H., Prost, R., and Friboulet, D. (2013). Pre-beamformed RF signal reconstruction in medical ultrasound using compressive sensing. *Ultrasonics*, 53(2):525–533.
- [Lilliefors1967] Lilliefors, H. W. (1967). On the Kolmogorov-Smirnov test for normality with mean and variance unknown. *Journal of the American statistical Association*, 62(318):399–402.
- [Lim and Nam2014] Lim, Y. and Nam, S. (2014). Target-to-clutter ratio enhancement of images in through-the-wall radar using a radiation pattern-based delayed-sum algorithm. *Journal of Electromagnetic engineering and Science*, 14(4):405–410.
- [Lin and Li2013] Lin, J. and Li, S. (2013). Nonuniform support recovery from noisy random measurements by orthogonal matching pursuit. *Journal of Approximation Theory*, 165(1):20–40.
- [Lopez et al.2018] Lopez, A., Bacelar, R., Pires, I., Santos, T. G., Sousa, J. P., and Quintino, L. (2018). Non-destructive testing application of radiography and ultrasound for wire and arc additive manufacturing. *Additive Manufacturing*, 21:298–306.
- [Love et al.2015] Love, L. J., Duty, C. E., Post, B. K., Lind, R. F., Lloyd, P. D., Kunc, V., Peter, W. H., and Blue, C. A. (2015). Breaking barriers in polymer additive manufacturing. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States). Manufacturing Demonstration Facility (MDF).
- [Lu and Wong2018] Lu, Q. Y. and Wong, C. H. (2018). Additive manufacturing process monitoring and control by non-destructive testing techniques: challenges and in-process monitoring. *Virtual and Physical Prototyping*, 13(2):39–48.
- [Lustig et al.2007] Lustig, M., Donoho, D., and Pauly, J. M. (2007). Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195.
- [Ma et al.2018] Ma, Y., Hong, H., and Zhu, X. (2018). Interaction multipath in through-the-wall radar imaging based on compressive sensing. *Sensors*, 18(2):549.

- [Malioutov et al.2010] Malioutov, D. M., Sanghavi, S. R., and Willsky, A. S. (2010). Sequential compressed sensing. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):435–444.
- [Mallat and Zhang1993] Mallat, S. and Zhang, Z. (1993). Matching pursuit with time-frequency dictionaries. Technical report, Courant Institute of Mathematical Sciences New York United States.
- [Martone et al.2010] Martone, A., Ranney, K., and Innocenti, R. (2010). Automatic through the wall detection of moving targets using low-frequency ultra-wideband radar. In *Radar Conference, 2010 IEEE*, pages 39–43. IEEE.
- [McCandless1999] McCandless, M. (1999). The MP3 revolution. *IEEE Intelligent Systems and their Applications*, 14(3):8–9.
- [Mohsin Riaz and Ghafoor2012] Mohsin Riaz, M. and Ghafoor, A. (2012). Through-wall image enhancement based on singular value decomposition. *International Journal of Antennas and Propagation*, 2012.
- [Mota et al.2015] Mota, J., Deligiannis, N., Sankaranarayanan, A. C., Cevher, V., and Rodrigues, M. (2015). Dynamic sparse state estimation using ℓ_1 - ℓ_1 minimization: Adaptive-rate measurement bounds, algorithms and applications. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 3332–3336. IEEE.
- [Mota et al.2014] Mota, J. F., Deligiannis, N., and Rodrigues, M. R. (2014). Compressed sensing with side information: Geometrical interpretation and performance bounds. In *Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on*, pages 512–516. IEEE.
- [Mota et al.2017] Mota, J. F., Deligiannis, N., and Rodrigues, M. R. (2017). Compressed sensing with prior information: Strategies, geometry, and bounds. *IEEE Transactions on Information Theory*, 63(7):4472–4496.
- [Muthukrishnan et al.2005] Muthukrishnan, S. et al. (2005). Data streams: Algorithms and applications. *Foundations and Trends® in Theoretical Computer Science*, 1(2):117–236.
- [Narayanan et al.2017] Narayanan, R. M., Gebhardt, E. T., and Broderick, S. P. (2017). Through-wall single and multiple target imaging using MIMO radar. *Electronics*, 6(4):70.
- [Needell and Tropp2009] Needell, D. and Tropp, J. A. (2009). CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and computational harmonic analysis*, 26(3):301–321.
- [Nkwari et al.2017] Nkwari, P., Sinha, S., and Ferreira, H. C. (2017). Through-the-wall radar imaging: A review. *IETE Technical Review*, pages 1–9.
- [Nyquist1928] Nyquist, H. (1928). Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644.
- [Oliver and Quegan2004] Oliver, C. and Quegan, S. (2004). *Understanding synthetic aperture radar images*. SciTech Publishing.
- [Ozdemir2012] Ozdemir, C. (2012). *Inverse synthetic aperture radar imaging with MATLAB algorithms*, volume 210. John Wiley & Sons.
- [Peabody Jr et al.2012] Peabody Jr, J., Charvat, G. L., Goodwin, J., and Tobias, M. (2012). Through-wall imaging radar. Technical report, Massachusetts Institute of Technology-Lincoln Laboratory Lexington United States.

- [Pourkamali-Anaraki and Becker2017] Pourkamali-Anaraki, F. and Becker, S. (2017). Pre-conditioned data sparsification for big data with applications to PCA and K-means. *IEEE Transactions on Information Theory*, 63(5):2954–2974.
- [Ribalta2011] Ribalta, A. (2011). Time-domain reconstruction algorithms for FMCW-SAR. *IEEE Geoscience and Remote Sensing Letters*, 8(3):396–400.
- [Rieder et al.2014] Rieder, H., Dillhöfer, A., Spies, M., Bamberg, J., and Hess, T. (2014). Online monitoring of additive manufacturing processes using ultrasound. In *11th European Conference on Non-Destructive Testing (ECNDT), Prague, Czech Republic, Oct.*, pages 6–10.
- [Rodriguez et al.2015] Rodriguez, E., Mireles, J., Terrazas, C. A., Espalin, D., Perez, M. A., and Wicker, R. B. (2015). Approximation of absolute surface temperature measurements of powder bed fusion additive manufacturing technology using in situ infrared thermography. *Additive Manufacturing*, 5:31–39.
- [Rudelson and Vershynin2005] Rudelson, M. and Vershynin, R. (2005). Geometric approach to error-correcting codes and reconstruction of signals. *International mathematics research notices*, 2005(64):4019–4041.
- [Schiffner and Schmitz2011] Schiffner, M. F. and Schmitz, G. (2011). Fast pulse-echo ultrasound imaging employing compressive sensing. In *Ultrasonics Symposium (IUS), 2011 IEEE International*, pages 688–691. IEEE.
- [Seifi et al.2016] Seifi, M., Salem, A., Beuth, J., Harrysson, O., and Lewandowski, J. J. (2016). Overview of materials qualification needs for metal additive manufacturing. *Jom*, 68(3):747–764.
- [Setlur et al.2011] Setlur, P., Amin, M., and Ahmad, F. (2011). Multipath model and exploitation in through-the-wall and urban radar sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 49(10):4021–4034.
- [Shannon1949] Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21.
- [Slotwinski2014] Slotwinski, J. A. (2014). Additive manufacturing: Overview and NDE challenges. In *AIP Conference Proceedings*, volume 1581, pages 1173–1177. AIP.
- [Stansbury and Idacavage2016] Stansbury, J. W. and Idacavage, M. J. (2016). 3D printing with polymers: Challenges among expanding options and opportunities. *Dental Materials*, 32(1):54–64.
- [Subotic et al.2008] Subotic, N., Keydel, E., Burns, J., Morgan, A., Cooper, K., Thelen, B., Wilson, B., Williams, W., McCarty, S., Lampe, B., et al. (2008). Parametric reconstruction of internal building structures via canonical scattering mechanisms. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 5189–5192. IEEE.
- [Sun et al.2014] Sun, X., Lu, B., and Zhou, Z. (2014). Moving target imaging and ghost mitigation in through-the-wall sensing application. *International Journal of Antennas and Propagation*, 2014.
- [Thanh et al.2008] Thanh, N., Van Kempen, L., Savelyev, T., Zhuge, X., Aftanas, M., Zaikov, E., Drutarovsky, M., and Sahli, H. (2008). Comparison of basic inversion techniques for through-wall imaging using UWB radar. In *Radar Conference, 2008. EuRAD 2008. European*, pages 140–143. IEEE.

- [Thompson et al.2016] Thompson, A., Maskery, I., and Leach, R. K. (2016). X-ray computed tomography for additive manufacturing: a review. *Measurement Science and Technology*, 27(7):072001.
- [Tivive et al.2011a] Tivive, F. H. C., Amin, M. G., and Bouzerdoum, A. (2011a). Wall clutter mitigation based on eigen-analysis in through-the-wall radar imaging. In *Digital Signal Processing (DSP), 2011 17th International Conference on*, pages 1–8. IEEE.
- [Tivive and Bouzerdoum2013] Tivive, F. H. C. and Bouzerdoum, A. (2013). An improved SVD-based wall clutter mitigation method for through-the-wall radar imaging. In *Signal Processing Advances in Wireless Communications (SPAWC), 2013 IEEE 14th Workshop on*, pages 430–434. IEEE.
- [Tivive et al.2011b] Tivive, F. H. C., Bouzerdoum, A., and Amin, M. G. (2011b). An SVD-based approach for mitigating wall reflections in through-the-wall radar imaging. In *Radar Conference (RADAR), 2011 IEEE*, pages 519–524. IEEE.
- [Tivive et al.2014] Tivive, F. H. C., Bouzerdoum, A., and Tang, V. H. (2014). Multi-stage compressed sensing and wall clutter mitigation for through-the-wall radar image formation. In *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2014 IEEE 8th*, pages 489–492. IEEE.
- [Trombetta et al.2017] Trombetta, R., Inzana, J. A., Schwarz, E. M., Kates, S. L., and Awad, H. A. (2017). 3D printing of calcium phosphate ceramics for bone tissue engineering and drug delivery. *Annals of biomedical engineering*, 45(1):23–44.
- [Tropp and Gilbert2007] Tropp, J. A. and Gilbert, A. C. (2007). Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on information theory*, 53(12):4655–4666.
- [Trzasko et al.2009] Trzasko, J., Haider, C., and Manduca, A. (2009). Practical nonconvex compressive sensing reconstruction of highly-accelerated 3D parallel MR angiograms. In *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*, pages 274–277. IEEE.
- [Van Den Berg and Friedlander2008] Van Den Berg, E. and Friedlander, M. P. (2008). Probing the Pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912.
- [Van Luong et al.2017] Van Luong, H., Deligiannis, N., Seiler, J., Forchhammer, S., and Kaup, A. (2017). Compressive online robust principal component analysis with multiple prior information. In *Signal and Information Processing (GlobalSIP), 2017 IEEE Global Conference on*, pages 1260–1264. IEEE.
- [Van Luong et al.2016] Van Luong, H., Seiler, J., Kaup, A., and Forchhammer, S. (2016). Sparse signal reconstruction with multiple side information using adaptive weights for multiview sources. In *Image Processing (ICIP), 2016 IEEE International Conference on*, pages 2534–2538. IEEE.
- [Wagner et al.2012] Wagner, N., Eldar, Y. C., and Friedman, Z. (2012). Compressed beamforming in ultrasound imaging. *IEEE Transactions on Signal Processing*, 60(9):4643–4657.
- [Walden1999] Walden, R. H. (1999). Analog-to-digital converter survey and analysis. *IEEE Journal on selected areas in communications*, 17(4):539–550.
- [Waller et al.2015] Waller, J. M., Saulsberry, R. L., Parker, B. H., Hodges, K. L., Burke, E. R., and Taminger, K. M. (2015). Summary of NDE of additive manufacturing efforts in NASA. In *AIP Conference Proceedings*, volume 1650, pages 51–62. AIP.

- [Wang et al.2012] Wang, W., Zhang, B., and Mu, J. (2012). Compressive SAR raw data with principal component analysis. *EURASIP Journal on Wireless Communications and Networking*, 2012(1):258.
- [Wang et al.2017] Wang, X., Jiang, M., Zhou, Z., Gou, J., and Hui, D. (2017). 3D printing of polymer matrix composites: A review and prospective. *Composites Part B: Engineering*, 110:442–458.
- [Wang and Fathy2012] Wang, Y. and Fathy, A. E. (2012). Advanced system level simulation platform for three-dimensional UWB through-wall imaging SAR using time-domain approach. *IEEE Transactions on Geoscience and Remote Sensing*, 50(5):1986–2000.
- [Ward2009] Ward, R. (2009). Compressed sensing with cross validation. *IEEE Transactions on Information Theory*, 55(12):5773–5782.
- [Wei et al.2010] Wei, S.-J., Zhang, X.-L., Shi, J., and Xiang, G. (2010). Sparse reconstruction for SAR imaging based on compressed sensing. *Progress In Electromagnetics Research*, 109:63–81.
- [Whittaker1915] Whittaker, E. T. (1915). Xviii.—on the functions which are represented by the expansions of the interpolation-theory. *Proceedings of the Royal Society of Edinburgh*, 35:181–194.
- [Wohlers et al.2016] Wohlers, T., Caffrey, T., and Campbell, I. (2016). Wohlers report 2016: 3D printing and additive manufacturing state of the industry annual worldwide progress report. Wohlers associates. Inc, Fort Collins, Colorado.
- [Xu and Xu2015] Xu, G. and Xu, Z. (2015). Compressed sensing matrices from Fourier matrices. *IEEE Transactions on Information Theory*, 61(1):469–478.
- [Xu et al.2012] Xu, Q., Yu, H., Mou, X., Zhang, L., Hsieh, J., and Wang, G. (2012). Low-dose X-ray CT reconstruction via dictionary learning. *IEEE transactions on medical imaging*, 31(9):1682–1697.
- [Yang et al.2015] Yang, D., Yang, X., Liao, G., and Zhu, S. (2015). Strong clutter suppression via RPCA in multichannel SAR/GMTI system. *IEEE geoscience and remote sensing letters*, 12(11):2237–2241.
- [Yang et al.2017] Yang, L., Hsu, K., Baughman, B., Godfrey, D., Medina, F., Menon, M., and Wiener, S. (2017). *Additive Manufacturing of Metals: The Technology, Materials, Design and Production*. Springer.
- [Yoon and Amin2008] Yoon, Y.-S. and Amin, M. G. (2008). Compressed sensing technique for high-resolution radar imaging. In *Signal Processing, Sensor Fusion, and Target Recognition XVII*, volume 6968, page 69681A. International Society for Optics and Photonics.
- [Yoon and Amin2009] Yoon, Y.-S. and Amin, M. G. (2009). Spatial filtering for wall-clutter mitigation in through-the-wall radar imaging. *IEEE Transactions on Geoscience and Remote Sensing*, 47(9):3192–3208.
- [Yu and Wang2009] Yu, H. and Wang, G. (2009). Compressed sensing based interior tomography. *Physics in medicine & biology*, 54(9):2791.
- [Zhang et al.2015] Zhang, S., Zhu, Y., Dong, G., and Kuang, G. (2015). Truncated SVD-based compressive sensing for downward-looking three-dimensional SAR imaging with uniform/nonuniform linear array. *IEEE Geoscience and Remote Sensing Letters*, 12(9):1853–1857.

- [Zhu and Wakin2015] Zhu, Z. and Wakin, M. B. (2015). Wall clutter mitigation and target detection using discrete prolate spheroidal sequences. In *Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa), 2015 3rd International Workshop on*, pages 41–45. IEEE.