SparseRI: A Compressed Sensing Framework for Aperture Synthesis Imaging in Radio Astronomy

Stephan Wenger
Marcus Magnor

Braunschweig : Computer Graphics Lab, 2010

(Technical Report 2010-1-11)

Veröffentlicht: 21.01.2010

http://www.digibib.tu-bs.de/?docid=00032047

Stephan Wenger
wenger@cg.tu-bs.de
Computer Graphics Lab, TU Braunschweig

Prof. Dr. Ing. Marcus Magnor
magnor@cg.tu-bs.de
Computer Graphics Lab, TU Braunschweig

SparseRI: A Compressed Sensing Framework for Aperture Synthesis Imaging in Radio Astronomy

Technical Report 2010-1-11

January 20, 2010

Computer Graphics Lab, TU Braunschweig

Contents

Abstract

In radio interferometry, information about a small region of the sky is obtained in the form of samples in the Fourier transform domain of the desired image. Since this sampling is usually incomplete, the missing information has to be reconstructed using additional assumptions about the image. The emerging field of Compressed Sensing (CS) provides a promising new approach to this type of problem which is based on the supposed sparsity of natural images in some transform domain. We present a versatile CS-based image reconstruction framework called *SparseRI*, an interesting alternative to the CLEAN algorithm, that permits a wide choice of different regularisers for interferometric image reconstruction. The performance of our method is evaluated on simulated data as well as on actual radio interferometry measurements from the VLA, showing that our algorithm is able to reproduce the main features of the test sources. The proposed method is a first step towards an alternative reconstruction approach that may be able to avoid typical artefacts like negative flux regions, work with large fields of view and non-coplanar baselines, avoid the gridding process, and reduce the amount of manual work that is required in order to obtain best-quality results.

1 Introduction

Since the middle of the twentieth century, interferometric techniques have been used to obtain images of the sky at radio wavelengths [RV46, RH60, SM68, PSB89]. The correlations between the signals from multiple antennae yield information about the frequency content of the image, eventually allowing to reconstruct the image itself. Since the size of the synthesised beam of such a telescope is inversely proportional to to the largest distance between any two antennae, very high spatial resolutions can be obtained this way. However, the process of reconstructing the image from incomplete frequency information is highly nontrivial, as missing information has to be appropriately reconstructed.

The reconstruction of missing information is only possible by specifying additional prior information about the image, such as its supposed smoothness, or the assumption that it contains a minimal amount of energy or a maximum amount of image 'entropy'. Traditionally, the iterative deconvolution algorithm CLEAN [Ḧ74] is used for the reconstruction of radio interferometric images. It implicitly assumes that the image is composed of a small number of point sources. Even though many extensions to the algorithm have been proposed (See Sect. 2), extended intensity distributions are still not always well reconstructed by the algorithm, and the process may require considerable user guidance in order to yield satisfactory results.

Besides, new telescopes like the Long Wavelength Array LWA [ECC$^+$09],

the Low Frequency Array LOFAR [vGN09] or the Square Kilometre Array SKA [Eke03, Sch04] that image large parts of the sky at once require reconstruction algorithms that handle increasing amounts of data as well as non-coplanar telescope geometries [CGB08, MS08]. Also, with a growing number of telescopes, a more automatic reconstruction pipeline is desirable.

A solution to many of these problems is proposed by the emerging theory of Compressed Sensing [CRT06a, CRT06b, Can06, Don06, Bar07]. It generalises and formalises the notion of 'prior information' about an image by relating it to the *compressibility* of the signal, a property that is present in most natural images [Cla85, DJL92] and distinguishes them from random signals. Compressed Sensing makes use of this fact in order to decide which of the possible images that explain a set of measurements is most probable, thereby turning the reconstruction problem into an optimization problem.

Today, a multitude of algorithms is available to efficiently solve the numerical problems that occur in Compressed Sensing [FNW07, TG07, WNF08], and the theory has successfully been applied to astronomy [BSO08], including radio interferometry [WJP+09, WPV09]. However, no Compressed Sensing reconstructions of real interferometric measurements have been published up to now. In this paper, we present SparseRI, a first step towards a general Compressed Sensing reconstruction framework for radio interferometry. We evaluate our algorithm using different sparsity priors on synthetic as well as real data. The evaluation shows that SparseRI is able to reproduce the main features of our test sources, and our reconstruction of real observations proves its practical applicability.

2 Interferometric image formation

In radio interferometry, correlations between multiple antennae are used to synthesise an aperture the size of the largest baseline. The measured *visibility* $V_\nu(\boldsymbol{b})$ for a baseline $\boldsymbol{b}$ depends on the intensity $I_\nu(\boldsymbol{s})$ of the sky in the direction $\boldsymbol{s}$ according to

$$V_\nu(\boldsymbol{b}) = \int I_\nu(\boldsymbol{s})\, \mathrm{e}^{-2\pi \mathrm{i} \boldsymbol{s}\cdot\boldsymbol{b}}\, \mathrm{d}\Omega. \tag{1}$$

Under the assumption that the sources being imaged are confined to a small region of the sky, this corresponds to a two-dimensional Fourier transform multiplied by a phase [PSB89]. This approximation may be rendered obsolete by Compressed Sensing techniques in the future (See Sect. 5).

Obtaining the sky image from the visibilities is an ill-posed inversion problem, and several approaches have been used to solve it. The CLEAN algorithm [H$\ddot{7}$4], the standard reconstruction method used in radio interferometry software like AIPS [1] and CASA [2], starts from the dirty image and

[1] http://aips.nrao.edu/

[2] http://casa.nrao.edu/

successively subtracts the point spread function around the brightest spots of the image. This process works well for images of isolated point sources, but does not always produce satisfying results for extended sources. Also, the convergence of the algorithm as well as the uniqueness of its solutions are not always guaranteed [Sch78, Sch79]. Regions that are supposed to be dark in the image are often excluded from the reconstruction process by hand in order to avoid artefacts in these regions, although some modifications to the algorithm try to reduce this and other objectionable effects [SF78, Cor83, Sch84] and to achieve better reconstructions of extended sources using a multi-scale approach [Cor08].

Other reconstruction algorithms used in the context of radio interferometry include the *maximum entropy method* [CE85], which attempts to maximise an entropy functional while satisfying the constraints imposed by the measured data, and more generally any algorithm that minimizes a plausibility function.

The Compressed Sensing methods that we are going to present allow for a mathematically very well-founded formulation of an important subclass of these algorithms. In simple terms, our algorithm searches the space of all images that explain the measurements for a solution that minimizes an appropriate metric, e.g. the magnitude of the coefficients of the image in a wavelet or gradient domain representation. Results similar to those of CLEAN can be obtained when the magnitude of the pixels is minimized.

3 Compressed Sensing

The theory of Compressed Sensing, recently introduced by Emmanuel Candès and others [CRT06a, CRT06b, Can06, Don06, Bar07], generalises the way of thinking about sampling. While the usual Shannon–Nyquist sampling theory requires that a band-limited signal be sampled with at least twice its highest frequency, Compressed Sensing states that a signal that is not necessarily band-limited but sparse in some basis (i.e. the majority of its coefficients in that basis are zero) can be reconstructed from a small number of measurements in another basis (i.e. a number of linear combinations of these coefficients). For example, if the shape of a signal is known to be Gaussian, it can be reconstructed perfectly from three known points even though the signal is not band-limited. Similarly, a signal that is sparse in the spatial domain (i.e. a combination of a small number of point sources) can be reconstructed from a small number of Fourier coefficients, which is typically the case in radio interferometry.

The measurement of a signal $\boldsymbol{x}$ using a linear measurement operator $\mathbf{M}$ can be written as $\boldsymbol{y} = \mathbf{M}\boldsymbol{x}$, where $\boldsymbol{y}$ is the resulting vector of measurements. In radio interferometry, $\boldsymbol{y}$ contains the visibilities, $\mathbf{M}$ is a Fourier transform, and $\boldsymbol{x}$ is the vector of image pixels. If $\boldsymbol{x}$ has a sparse representation $\boldsymbol{s} = \mathbf{B}\boldsymbol{x}$

in a sparsity basis $\mathbf{B}$, this turns into $\boldsymbol{y} = \mathbf{M}\mathbf{B}^{-1}\boldsymbol{s} = \mathbf{A}\boldsymbol{s}$, with $\mathbf{A} = \mathbf{M}\mathbf{B}^{-1}$.

The possibility of perfect reconstruction of the signal $\boldsymbol{x}$ depends on the properties of the combined measurement and sparsity matrix $\mathbf{A}$. Candès and Tao [CT05] prove that perfect reconstruction occurs when (in simple terms) all sets of a sufficiently high number of columns of $\mathbf{A}$ are approximately orthonormal. The required number of columns depends on the sparsity level of the signal. In particular, any *random* measurement satisfies this criterion with high probability [Can06], so that in this sense a random Fourier sampling contains more information than a quasi-regular baseline pattern. Stable reconstruction is possible even when $\boldsymbol{y}$ is perturbed by measurement error and when $\boldsymbol{s}$ is not perfectly sparse [CRT06b], two properties that are important for real-world applications.

In order to reconstruct a sparse signal $\boldsymbol{s}$ from its measurement vector $\boldsymbol{y}$, one has to find the sparsest vector $\boldsymbol{s}$ that satisfies $\boldsymbol{y} = \mathbf{A}\boldsymbol{s}$ [CRT06a]. Finding a sparsest vector is equivalent to minimizing the ℓ_0 norm which is defined as the number of nonzero coefficients of a vector. Unfortunately, the ℓ_0 norm minimization problem is computationally not tractable. With increasing dimension of the vector $\boldsymbol{s}$, however, the solution of the ℓ_0 norm minimization approaches that of an ℓ_1 norm minimization which can be efficiently solved. Algorithms also exist when the entries of $\boldsymbol{y}$ and $\boldsymbol{s}$ are vectors, for example different polarisation components of a signal [FR08].

As a non-linear optimization approach, Compressed Sensing is potentially more powerful than conventional linear optimization or deconvolution approaches. The optimization result is also mathematically uniquely defined, as opposed to the procedural definition of CLEAN. Therefore it does not depend on any specific implementation or parameter set, and theoretical analysis of the algorithm is greatly simplified. The only parameters that influence the optimal solution are the expected noise level of the measurement and the chosen sparsity prior.

3.1 Sparsity priors

The selection of a sparsity basis or *sparsity prior* represents our assumptions about the image we want to reconstruct and as such can strongly influence the reconstruction. For example, the pixel basis is obviously well suited to represent the assumption of isolated point sources. In contrast, most terrestrial images are likely to contain large regions of homogeneous or slowly changing intensity, possibly with small-scale perturbations or sharp edges, and appropriate sparsity bases are known which can as well be used for astronomical imaging. For example, different wavelet representations like those proposed by Daubechies [Dau88] or Cohen et al. [CDF92] efficiently compress many natural images because they provide a scaling-independent but localised basis. Minimising the total variation (which is not a basis in the strict sense) will reliably localise extended emissive regions because

sharp edges are not penalised more than smooth gradients in this case. In Sect. 4, we provide sample reconstructions using pixel and total variation priors.

3.2 Reconstruction algorithm

In previous work [WJP+09, WPV09], Compressed Sensing techniques were used to recover simulated images of different radio sources (random Gaussians and string signals) from simulated measurements as a proof of concept. However, the algorithm has not been tested on actual interferometric measurements. We implement a different algorithm that can be shown to work with data from real radio interferometers like the Very Large Array (VLA) and that permits a wide choice of different regularisers in order to reconstruct a wider range of sources. It is also stable with respect to the errors introduced by the physical measurement process and as such can be used on real interferometric observations.

Our algorithm is an adaptation of the very general Compressed Sensing framework of Wright et al [WNF08]. The algorithm minimizes a functional of the form $||\boldsymbol{y} - F(\boldsymbol{x})||_2^2 + \lambda f(\boldsymbol{x})$, where $\boldsymbol{y}$ is the vector of measurements, F is the Fourier transform and $\boldsymbol{x}$ is the vector of image pixels to be reconstructed. The first term (the 'data term') promotes $\boldsymbol{x}$ to be consistent with the measurements, while the second term represents the plausibility metric. The parameter λ can be chosen so that $||\boldsymbol{y} - F(\boldsymbol{x})||_2^2$ lies within the expected noise level. The plausibility metric $f(\boldsymbol{x})$, in our case, is either the ℓ_1 norm (defined as $(\sum_i |x_i|^p)^{(1/p)}$) of the coefficients of $\boldsymbol{x}$ in some sparsity basis, or the total variation [CD09], the ℓ_1 norm of the gradient image of $\boldsymbol{x}$.

The Fourier transform is implemented as a two-dimensional Fast Fourier Transform (FFT) on a regular grid for efficiency reasons. The necessary *gridding* process inevitably introduces quantisation errors. Unfortunately, convolution-based resampling cannot directly be used to reduce this error because it complicates the computation of the data term, but gridding error may be avoided entirely by future extensions of our framework (See Sect. 5).

4 Evaluation

In the following, we evaluate the performance of SparseRI using two different datasets, a synthetic one and one from the Very Large Array.

4.1 Simulated data

For a simulated source, the 'ground truth' image is known, allowing us to objectively compare the performance of different reconstruction algorithms. We provide three different error metrics: the signal-to-noise ratio (SNR),

which is defined as $\mathrm{SNR} = -20 \log_{10} \frac{\sigma_{\mathrm{residual}}}{\sigma_{\mathrm{original}}}$ [WJP+09], the RMS error normalised to the average of the true image, and the dynamic range, which is the ratio of the highest peak in the reconstruction to the standard deviation of the background noise. Since SparseRI and CLEAN are non-linear reconstruction algorithms, the dynamic range has to be interpreted with care.

In Fig. 1, we present two simulated sources (b) that are reasonably close to real extended radio sources while not being particularly easy or difficult to reconstruct for either algorithm: a uniform gradient with Gaussian decay and a series of Gaussians of increasing size. The measurements were simulated using the CASA SIMDATA task. The fairly high amount of emission, together with comparatively low sampling density (UV coverage (a) is 4.2 per cent), pushes the algorithms to their limits so that the differences in reconstruction quality become visible.

As a reference, CLEAN reconstructions (c) were made using Clark's algorithm [Cla80] algorithm from the CASA software package. In order to make the results user-independent and comparable, both reconstruction algorithms have been run with their default parameters (for the series of Gaussians, however, the number of CLEAN iterations had to be increased to 90000 in order to yield satisfying results). For the CLEAN reconstructions, the SNRs are 6.1 dB and 15.9 dB, the dynamic ranges are 17.6 and 121.5, and the RMS errors per pixel are 0.12 and 0.05 per cent of the true mean intensity for the first and second source, respectively. It is noticeable that in the second source, the Gaussians are not clearly separated from each other.

The SparseRI reconstructions (d) yield SNRs of 5.0 dB and 6.6 dB, dynamic ranges of 19.2 and 24.6, and RMS errors per pixel of 0.13 and 0.15 per cent of the true mean intensity. These metrics indicate that the performance of SparseRI is at least comparable with CLEAN, although the emerging algorithm is not yet on par with its mature counterpart. The largest contribution to the lower SNR values is due to the occurrence of negative flux regions which are not penalized in the current implementation. However, visual comparison shows that SparseRI is able to resolve the series of Gaussians, even though some larger scale stripes are present in the background. The computation time of the non-optimized, single-threaded SparseRI algorithm was about 4.5 seconds on conventional PC hardware. The automatic CLEAN reconstruction took three seconds for the 1000 iterations of the first source, and 160 seconds for the 90000 iterations of the second source.

4.2 Real data

The applicability of our algorithm to actual noise-affected measurements is demonstrated using snapshot observations from the VLA in the 'D' configuration at 14.965 GHz and a UV coverage of 2.4 per cent. Fig. 2 shows a

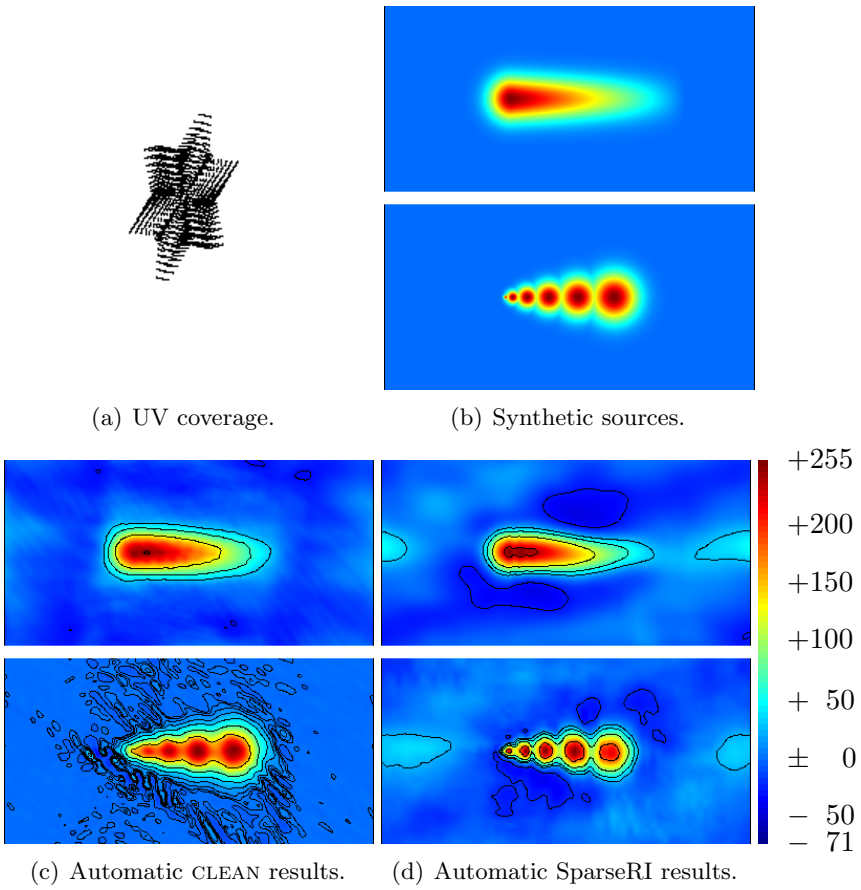

(a) UV coverage. (b) Synthetic sources.

(c) Automatic CLEAN results. (d) Automatic SparseRI results.

Figure 1: Synthetic sources (b) with reconstructions from simulated measurements using CLEAN with the default parameters (c) and SparseRI's Daubechies wavelet minimization (d), respectively. Contour levels are plotted – except for the noise-free originals (b) – at $-1, 1, 2, 4, \ldots, 2^n$ times the respective 3σ RMS noise. All images are 256 by 256 pixels, and the UV coverage (a) is 4.2 per cent.

series of reconstructions of this dataset containing the Sgr A West region, including the central 'Minispiral'. Fig. 2(b) shows the reconstruction using CLEAN in CASA with default parameters. However, most astronomers will not use default parameters in AIPS or CASA, and in Fig. 2(d) we therefore also include a map cleaned more traditionally using AIPS (done by co-author YP). Finally, Fig. 2(c) shows the results from SparseRI using the pixel basis. The SparseRI map is convolved with the CLEAN beam so that spatial resolution becomes comparable. SparseRI appears to better reconstruct the image than the automated CLEAN, but still shows some systematic imaging effects at lower intensity levels. The dynamic range of the reconstructions (assuming the right quarter of the image is background) is 43.5 for CLEAN with the default parameters, 1367 for user-guided CLEAN and 610 for SparseRI. The computation time of the non-optimized, single-threaded SparseRI algorithm was 9 seconds on conventional PC hardware, the automatic CLEAN reconstruction took about two seconds, while approximately 15 minutes were needed for the user-guided reconstruction, including self-calibration steps.

5 Conclusion

The evaluation of our algorithm on simulated and real measurements shows that SparseRI is able to provide interferometric image reconstructions that reproduce the main features of complex sources without any manual parameter tweaking or boxing, at comparable computation times as traditional reconstruction algorithms. While SparseRI is a non-linear algorithm, it has been proven to converge towards the optimal solution as well as to be stable with respect to noise [CRT06b]. The reconstruction results of this first implementation are still above the observational noise level. However, there exist a number of beneficial constraints and other potential improvements that we have not yet explored. We are confident that the SparseRI approach has the potential to become on par and even surpass traditional reconstruction algorithms with respect to automated performance and achievable resolution.

We plan to systematically evaluate different sparsity bases in a follow-up paper, which may be a first step towards a sparsity basis that is explicitly designed for typical radio images. Such a basis might be found by statistically analysing a large database of such images. However, one would have to make sure that the database does not suffer from systematic errors (such as reconstruction artefacts in the training database), and an efficient representation for such a non-systematic basis has yet to be found. A combination of the previously discussed wavelet, pixel and total variation representations would also be able to increase the reconstruction quality on a wider range of sources. In general, the choice of sparsity basis will reflect which features or spatial frequencies the user wants to emphasize.

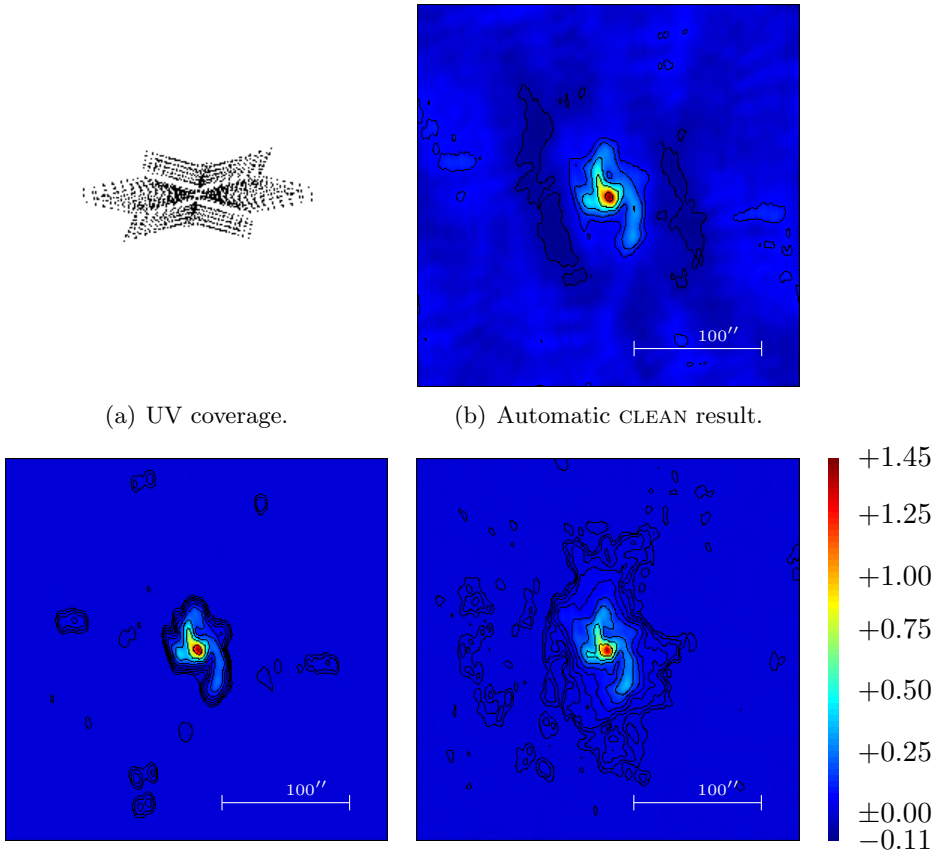

(a) UV coverage. (b) Automatic CLEAN result.

(c) Automatic SparseRI result. (d) User-guided CLEAN result.

Figure 2: Sgr A West reconstructed from VLA data using CLEAN with the default parameters (b), SparseRI's pixel magnitude minimization (c), and CLEAN with user guidance (d). Contour levels are plotted at $-1, 1, 2, 4, \ldots, 2^n$ times the respective 3σ RMS noise. The color scale is in arbitrary units, with the peak flux of all images normalized to the same level. All images are 256 by 256 pixels (300 by 300 seconds of arc), and the UV coverage (a) is 2.4 per cent.

Incorporating further constraints into SparseRI – such as suppressing negative flux regions, which cause most of the reconstruction error in the current implementation – will further increase reconstruction quality [WJP+09]. Additionally, we aim to include calibration with respect to an absolute flux scale, which is not yet implemented in the present version.

We also consider implementing a fast, parallelised three-dimensional non-uniform Fourier transform, for example on inexpensive graphics hardware. Such a general transform could prevent the hitherto unavoidable gridding error, possibly further decrease computation time, and enable fully three-dimensional aperture synthesis for recent interferometry arrays with large fields of view and very long (non-coplanar) baselines, such as the EVN, the VLBA, or the recent LWA, LOFAR and SKA systems.

Acknowledgments

The authors would like to thank Sanjay Bhatnagar, Soheil Darabi, Karl-Heinz Glaßmeier, Dave Martin, Ylva Pihlström, Urvashi Rau and Pradeep Sen for their support and advice.

References

[Bar07] R. Baraniuk. Compressive sensing. *IEEE Signal Proc. Magazine*, 24:118–121, 2007.

[BSO08] J. Bobin, J.-L. Starck, and R. Ottensamer. Compressed sensing in astronomy. *IEEE J. Selected Topics in Signal Proc.*, 2:718–726, 2008.

[Can06] E. J. Candès. Compressive sampling. *Int. Congress Math.*, 3:1433–1452, 2006.

[CD09] A. Chambolle and J. Darbon. On total variation minimization and surface evolution using parametric maximum flows. *Int. J. Computer Vision*, 84:288–307, 2009.

[CDF92] A. Cohen, I. Daubechies, and J. C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Comm. Pure and Appl. Math.*, 45:485–560, 1992.

[CE85] T. J. Cornwell and K. F. Evans. A simple maximum entropy deconvolution algorithm. *A&A*, 143:77–83, 1985.

[CGB08] T. J. Cornwell, K. Golap, and S. Bhatnagar. The noncoplanar baselines effect in radio interferometry: The W-projection algorithm. *IEEE J. Selected Topics in Signal Proc.*, 2:647–657, 2008.

[Cla80] B. G. Clark. An efficient implementation of the algorithm 'CLEAN'. *A&A*, 89:377–378, 1980.

[Cla85] R. J. Clarke. *Transform coding of images*. Academic Press, 1985.

[Cor83] T. J. Cornwell. A method of stabilizing the CLEAN algorithm. *A&A*, 121:281–285, 1983.

[Cor08] T. J. Cornwell. Multiscale CLEAN deconvolution of radio synthesis images. *IEEE J. Selected Topics in Signal Proc.*, 2:793–801, 2008.

[CRT06a] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52:489–509, 2006.

[CRT06b] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure and Appl. Math.*, 59:1207–1223, 2006.

[CT05] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51:4203–4215, 2005.

[Dau88] Ingrid Daubechies. Orthonormal bases of compactly supported wavelets. *Comm. Pure and Appl. Math.*, 41:909–996, 1988.

[DJL92] R. A. DeVore, B. Jawerth, and B. J. Lucier. Image compression through wavelet transform coding. *IEEE Trans. Inform. Theory*, 38:719–746, 1992.

[Don06] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52:1289–1306, 2006.

[ECC+09] S. W. Ellingson, T. E. Clarke, A. Cohen, J. Craig, N. E. Kassim, Y. Pihlstrom, L. J. Rickard, and G. B. Taylor. The Long Wavelength Array. *Proc. IEEE*, 97:1421–1430, 2009.

[Eke03] R. D. Ekers. Square Kilometre Array (SKA). In *Proc. of the IAU 8th Asian-Pacific Regional Meeting*, volume 289, pages 21–28, 2003.

[FNW07] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE J. Selected Topics in Signal Proc.*, 1:586–598, 2007.

[FR08] M. Fornasier and H. Rauhut. Recovery algorithms for vector-valued data with joint sparsity constraints. *SIAM J. Num. Anal.*, 46:577–613, 2008.

[H$\ddot{7}$4] J. A. Högbom. Aperture synthesis with a non-regular distribution of interferometer baselines. *A&AS*, 15:417–426, 1974.

[MS08] J. D. McEwen and A. M. M. Scaife. Simulating full-sky interferometric observations. *MNRAS*, 389:1163–1178, 2008.

[PSB89] R. A. Perley, F. R. Schwab, and A. H. Bridle, editors. *Synthesis Imaging in Radio Astronomy*, volume 6 of *Conference Series*. Astronomical Society of the Pacific, 1989.

[RH60] M. Ryle and A. Hewish. The synthesis of large radio telescopes. *MNRAS*, 120:220–230, 1960.

[RV46] M. Ryle and D. D. Vonberg. Solar radiation on 175 Mc./s. *Nat*, 158:339–340, 1946.

[Sch78] U. J. Schwarz. Mathematical-statistical description of the iterative beam removing technique (method CLEAN). *A&A*, 65:345–356, 1978.

[Sch79] U. J. Schwarz. The method CLEAN – use, misuse and variations. In *IAU Colloq. 49: Image Formation from Coherence Functions in Astronomy*, volume 76, pages 261–275, 1979.

[Sch84] F. R. Schwab. Relaxing the isoplanatism assumption in self-calibration; applications to low-frequency radio interferometry. *AJ*, 89:1076–1081, 1984.

[Sch04] R. T. Schilizzi. The Square Kilometer Array. In *SPIE Conference Series*, volume 5489, pages 62–71, 2004.

[SF78] A. Segalovitz and B. R. Frieden. A 'CLEAN'-type deconvolution algorithm. *A&A*, 70:335–343, 1978.

[SM68] G. W. Swenson Jr. and N. C. Mathur. The interferometer in radio astronomy. *Proc. IEEE*, 56:2114–2130, 1968.

[TG07] J. Tropp and A. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Inform. Theory*, 53:4655–4666, 2007.

[vGN09] M. de Vos, A. W. Gunst, and R. Nijboer. The LOFAR telescope: System architecture and signal processing. *Proc. IEEE*, 97:1431–1437, 2009.

[WJP+09] Y. Wiaux, L. Jacques, G. Puy, A. M. M. Scaife, and P. Vandergheynst. Compressed sensing imaging techniques for radio interferometry. *MNRAS*, 395:1733–1742, 2009.

[WNF08] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo. Sparse reconstruction by separable approximation. In *IEEE Int. Conf. Acoustics, Speech and Signal Proc.*, pages 3373–3376, 2008.

[WPV09] Y. Wiaux, G. Puy, and P. Vandergheynst. Compressed sensing reconstruction of a string signal from interferometric observations of the cosmic microwave background. *MNRAS*, 2009. in press.