The Maxwell Advanced Technology Fund

2010 – 2015

THE UNIVERSITY *of* EDINBURGH

CONTENTS

THE MAXWELL ADVANCED TECHNOLOGY FUND

INTRODUCTION

Dear Sponsor,

The Maxwell Advanced Technology Fund

I would like to thank you once again for your continued generosity in supporting postgraduate students within the Institute for Digital Communications and to take this opportunity you update you on the exciting research that has been generated from your generous donations so far.

Currently the Maxwell Advanced Technology Fund has helped three PhD students: Chunli Guo, Ashley Hughes and Jonathan Mason.

Chunli's research was on the emerging field of compressed sensing – a new theory on signal and image reconstruction from partial measurements. Currently the University of Edinburgh is researching a number of applications for compressed sensing including: Radar, MRI and CT imaging. Chunli's research applied concepts from information theory to develop new reconstruction algorithms for compressed sensing leading in one case to a 20 times speed up in computation over previous state-of-the art techniques.

Chunli successfully defended her thesis on 10/12/2014 and been awarded a post-doctoral position at University College London, working on an MOD funded project on processing for sonar imaging. Her research generated 7 publications (3 IEEE journal papers and 4 international conference papers) and she was also invited to present her research at an event for advanced CT imaging in Denmark. Another output from her research was a new collaboration with Prof. Goertz at the Technical University of Vienna (TUV) on the links between information theory and compressed sensing (leading to one of the journal publications). IDCOM will subsequently receive a visiting student from TUV next spring to continue this work. Ideas from her work are also being used now in collaboration with GE Global Research working on advanced image formation for CT baggage scanners to provide improved explosives detection. This work is funded by the US Dept. Homeland Security.

Ashley's thesis has been in the area of audio signal processing, and addresses Acoustic Source Localisation and Tracking (ASLT) using microphone arrays. Acoustic tracking has application in, for example, security, human-computer interaction, and audio diarisation of meetings. Ashley's thesis specifically addresses improving significantly the computational performance of ASLT techniques using a variety of optimisation techniques.

Ashley submitted his thesis in June 2015, and recently successfully defended his thesis on 2nd October. He has published two conference papers covering his work, one of which focused on audio-visual tracking and linked with our EPSRC/DSTL funded defence signal processing research (University Defence Research Collaboration). The computationally efficient algorithms presented in Ashley's thesis also form an important component for a proposed EPSRC project, in collaboration with Informatics at the University of Edinburgh, on self-localising ad-hoc microphone arrays for applications such as ambient assisted living. Fast ASLT algorithms will be crucial for the real-time implementation, as well as reducing power consumption of the proposed sensor network.

Since January of 2015, Ashley has been working full-time in a graduate position within the DSP software team at Dialog Semiconductors in Edinburgh. His work in Dialog is within the area of digital audio for low-powered devices, and will no doubt benefit hugely from his expertise obtained during his graduate studies.

Our latest Maxwell Advanced Technology Fund recruit, Jonathan, is just coming up to the end of the first year of his PhD studies and is investigating how advanced image processing can be used in image guided radiotherapy. The goal is to be able to use previously taken CT images to enhance the imaging performed during radiotherapy opening up the possibility of adaptive radiotherapy where a radiotherapy plan is modified based on near real-time image data generated on the day of treatment. Jonathan's preliminary results look very promising and I hope we will be able to present these in a conference paper soon. The funding of Jonathan Mason has enabled IDCOM to build stronger links with Dr Nailon in Oncology at the Western General Hospital, and has helped lead to further collaborative work, including the secondment of Dr Laurenson from IDCOM next year to work on automated tumour detection.

The Maxwell Advanced Technology Fund has provided IDCOM with an incredibly valuable resource that has enabled us to be more flexible in our recruitment and allowed us to offer funded places to students outside of the usual EPSRC Doctoral Training timeline. This seems to have had the added benefit of helping us secure more UK students: 2 out of the 3 so far.

Thank you once again.

Mike Davies CEng, MIET, FIEEE
Professor of Signal and Image Processing
Head of Institute for Digital Communications & the Joint Research Institute for Signal and Image Processing
University of Edinburgh

THE MAXWELL ADVANCED TECHNOLOGY FUND

SCHOLARS

Details of the research undertaken by each student supported by the Fund follows

Jonathan Mason
Engineering (Digital Communications) PhD

Research Project and its Context

Background

The area of investigation for the research project is in the use of image guidance in radiation therapy for cancer. Specifically, it looks at how the on board cone beam computed tomography (CBCT) units on linear accelerators can give an accurate image of the biological tissues of interest, from which their shape and location can be precisely determined. From this information, it is hoped to be able to then adjust the treatment plan to more effectively irradiate the cancer, whilst sparing healthy tissue. Two large difficulties in attempting to use this technique with standard protocols is that the accuracy of reconstructed images in rather poor, and there is a large associated dose with acquiring so many measurements with an ionising radiation. Fortunately, in this setting, there is a wealth of information that could hopefully be exploited to be able to generate a high fidelity image from only a small number of measurements. Specifically, we believe that by utilising the planning CT image, this should be possible, and this serves as the main theme of the project.

One powerful technique for finding high quality solutions from a system with few measurements is that of compressed sensing, which in the context of image reconstruction, usually regularises some representation of this image to be sparse. In our setting, we wish to extend this to regularise some representation of the image with respect to our prior planning scan to be sparse. This process is known as prior image constrained compressed sensing (PICCS) [1], and will serve as a prototype for the project.

"This would not have been possible without the support of the scholarship, so I am truly grateful to be the recipient."

Progress to Date

The progress so far has been managing to generate a suitable framework to develop in, implementing a number of different algorithms to perform PICCS, and to applying these both to synthetic data and to those generated from a real scanner. Initially, we started developing in a pre-existing toolbox written in C++ known as RTK, which includes a decent number of algorithms specifically for CBCT, and has already seen application of compressed sensing techniques [2]. Unfortunately, it proved to be rather inflexible to work in, and was written with such abstraction that it was difficult to write even the simplest of operations. Due to this, we made the decision to switch to Mat lab, where algorithmic development was both rapid and powerful.

Although there already exists a toolbox available for CBCT reconstruction [3], this was unable to operate directly on raw data from a scanner, and it proved very slow for large systems. Instead, we made our own functions to perform reconstruction from CBCT data, and had it able to also work with real scanner data. With just the operations of back-projection and forward-projection, we can now implement any reconstruction algorithm, and have the potential to use many well established toolboxes for compressed sensing reconstruction on both synthetic and real data.

As a simple demonstration of the power of our operators, we made an implementation of the popular FDK [4] technique, and compared it against the implementation in [3] and against two versions of the RTK toolbox for the case of a 128 ×128 ×128 voxel Sheep-Logan Phantom scanned with 360 cone beam projections onto a square detector of 128 × 128 pixels. The results of this test are shown in Table 1, which indicates

The speed and accuracy of our implementation. The reason for the substantial speed advantage of the RTK method is a case of a compiled language against an interpreted one. The good accuracy of our method however does compensate for this to some extent. We think the reason for the poor performance of RTK is due to the limited filter choice, which we were able to optimise for the Mat lab codes. This result is further illustrated in Figure 1.

	Time (s)	PSNR (dB)
Kim	128.4	45.9
Mason	17.4	46.7
RTK 1.0	8.4	41.7
RTK 1.1	3.4	41.7

Table 1: Performance of FDK algorithm: **Kim** is the alternative **Mat lab** CBCT toolbox by [3]; **Mason** is our implementation; **RTK** 1.0 is from the stable release of the RTK toolbox; **RTK** 1.1 is the development version of the RTK toolbox with several accelerations.

(a) Specimen (b) FDK Reconstruction (c) RTK Reconstruction

Figure 1: Central slices of test phantom: a is the original scanned phantom; b is the reconstruction using an FDK implementation with our operators; c is the FDK reconstruction using the RTK toolbox. The intensities have been scaled to improve contrast.

Along with being able to operate on simulated data like the Sheep-Logan phantom, our operators can also work on data from a real CBCT system. In Figure 2, we show two slices from a volume of $512 \times 512 \times 4$ voxels of a physical phantom scanned using a Varian on-board imaging device (OBI). This was scanned from 49 cone beam projections onto a flat detector of 1024×768 pixels, and reconstructed using our FDK implementation in ~ 3 seconds. The streaking artefacts seen are very characteristic of the FDK method for a limited number of views.

After producing a framework in Mat lab for the CBCT system, we next started investigating several algorithms for performing PICCS. To do this, we focused on so-called proximal algorithms [5], which are well suited for solving problems with separable objectives. In our case, the two objectives are forming a reconstruction that is a close match to the data, and having a regularised solution against our prior planning CT image. Throughout this work, we attempted to solve the following objective function

$$\text{minimize} \quad \frac{1}{2\lambda}(Ax - b)^T \Sigma (Ax - b) + \alpha \|x\|_{TV} + (1-\alpha)\|x - \tilde{x}\|_{TV}, \qquad (1)$$

where A is the matrix describing the CBCT system, b is a column vector containing all the measurements, x is the reconstructed volume represented also as a column vector, $\tilde{x}$ is the planning CT, and λ and α are scalar constants to balance the three terms in the expression. The first term in Equation 1 describes the weighted error between the reconstruction and the measurements, and the diagonal matrix σ has values related to the magnitude of each measurement to account for their Poisson distributed nature. The other two terms are both total variation (TV) semi-norms, which corresponds to the sum of the discrete gradients. Minimising the first TV term promotes an image with a sparse discrete gradient, and minimising the second one reduces the TV of the difference between this and the planning CT. Finding a solution to Equation 1 is a very similar formulation to [1]. Whether this is truly a good thing to do, will likely be determined from experimentation on real data, and assessment from a clinician.

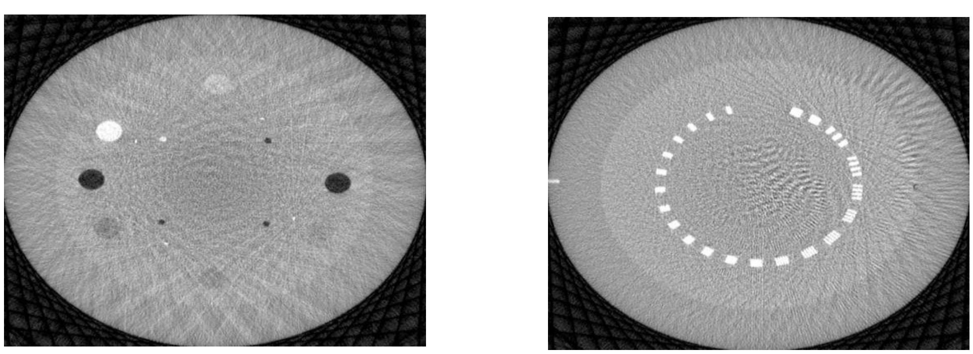

(a) Slice 1 (b) Slice 2

Figure 2: Two slices from the FDK reconstruction from 49 projections of the clock phantom scanned using a Varian OBI: a is a slice from a section designed to test contrast; b is a slice from a different section designed to test resolution.

In order to solve the objective above, we proposed the use of a number of algorithms. A preliminary set of results from these methods applied to the simple Sheep-Logan phantom is plotted in Figure 3.

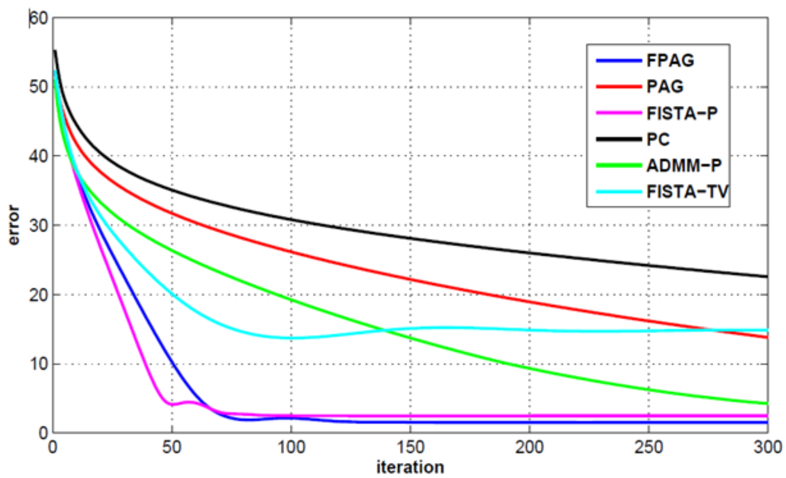

Figure 3: Progression of the error for various proximal methods. It should be noted that the step size for the ADMM based methods (ADMM-P and PC) was found to be safe but slow.

All the algorithms in Figure 3 solve various versions of Equation 1. All the methods are proximal methods: the proximal average gradient (PAG), fast PAG (FPAG), prior constrained FISTA [6] (FISTA-P) and to constrained FISTA [7] (FISTA-TV) are based upon the proximal gradient descent; and the proximal consensus (PC) and prior constrained ADMM (ADMM-P) are based upon the alternated direction method of multipliers (ADMM) [5]. We implemented each of them to use our operators, and they all resulted in reconstructions of pleasing quality.

Expected Contribution to Knowledge and Impact

Image guided radiotherapy could have the potential to significantly increase the efficiency of dose delivery, and hopefully will mean a reduced recurrence of cancer and reducing the treatment's toxicity to healthy tissue. For this reason, the techniques developed throughout this project may have a large positive impact. Although similar efforts to utilise planning CT information are currently being explored by others [8][9], they are still a long way off clinical practice and there are still many open questions in generating an optimal reconstruction for this application. We therefore expect to contribute knowledge about the degree to which prior information can be exploited in this setting, and also present a solution for doing this that works well in practice. In order to increase the impact of our work, we plan to make publicly available both our developed software and data, so that others can use our methods and also extend them easily.

Programme and Methodology

Aims and Objectives

To begin with, the aims of the project will be to develop an algorithm for CBCT reconstruction using planning CT data, to generate a highly accurate volume in a short amount of time. The set objectives for the end of the project, in no particular order, are as follows:

1. Generate images deemed acceptable by clinical radiographers;
2. Develop a method that could take no longer that 1 minute to run on some hardware;
3. be able to generate treatment plans on the images, and have this also deemed acceptable by radiographers;
4. Require a level of radiation considered 'low' by clinical radiographers;
5. Have a robust method that is proved to converge for all its intended uses;
6. Establish an optimal parameter setting free of ambiguity;
7. Determine the optimum utilisation of a given dose for the best image.

Methodology

We believe the key to achieving many of the objectives is through interaction. In the case of 1, 3 and 4, this naturally will be done by talking with clinical radiographers, showing them various results, and establishing what it is they deem acceptable. It is also likely that items 5 and 6 could be achieved through extensive experimentation and also having discussions with experts in mathematical optimisation.

At this point, we make the prediction that the objective can be achieved through extension of the PICCS methods already implemented. The methodology for doing this will be in getting an algorithm that is fast and either provably stable or empirically so with very high confidence, then apply it to real data and get clinicians to assess its quality. We think that likely contenders for this are in the approximate message passing (AMP) [10] or in stochastic gradient descent [11], and will therefore investigate both of these for the application. Two more avenues of research will be in using regularisations that could be more suitable than TV, and in how best to register the planning CT onto the CBCT measurements. We would like to speculate that it may well be the case that the solution to the problem is just a registration operation, and we are open to this being the case.

Practically, we plan to go about this method by continuing the recent development efforts so far, by first implementing and testing algorithms for synthetic data in Matlab, then applying them to real data. After this point, we propose to get opinions form clinicians, then taking this feedback to adjust the algorithms, and continue doing this in a loop of development.

Project Management

The management of the project will be maintained by generating regular reports on its progress and getting the supervisors and peers to review them. If we can keep a representative log of the work in such a way, then it is hoped that the project will be able to stay managed. Additionally, by having produced this work, it should be easier to then publish when contributions start to arise.

"I feel privileged to be able to work full time on something that I carry such an interest for, and being given the opportunity to work with such skilled colleagues within state of the art facilities is fantastic."

With reference to the diagrammatic flow chart at the end of this report, we intend to spend a decent amount of time doing research. Although this research has been allocated to specific areas at this point, there is bound to be new directions that result from reading around and those turn in main themes. Nevertheless, what is presented is a best guess at where the time will be spent over the rest of the project. The strong distinctions between activities is not intended to be strict. In fact, it is likely that all of the listed items and more will be investigated or considered in any given time period. Instead, its use is in ensuring a development and publication cycle is established, and if we find any strong deviations from this, then we may need to address it.

Resources

There are a few critical resources for this project, but fortunately they seem to be well placed at this moment. Since the contract with the NHS has now been set up, access to physical scanners and the data they produce should be possible shortly. On top of this, we have access now to a very powerful computer location at the Western General Hospital that should allow our algorithms to run at a more representative speed. This is important in addressing item 2 in the objectives.

References

[1] Gang-Hong Chen, Jie Tang, and Shuai Leng. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly under sampled projection data sets. Medical physics, 35(2):660–663, 2008.

[2] Cyril Mory, Vincent Auvray, Bo Zhang, Michael Grass, Dirk Schafer, S James Chen, John D Carroll, Simon Rit, Fran¸coise Peyrin, Philippe Douek, and Lo¨ıc Boussel. Cardiac C-arm computed tomography using a 3D + time ROI reconstruction method with spatial and temporal regularization. Medical physics, 41:021903, 2014.

[3] Kyungsang Kim. MATLAB CBCT Toolbox, 2015.

[4] L. a. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. Journal of the Optical Society of America A, 1(6):612, 1984.

[5] Neal Parikh and Stephen Boyd. Proximal Algorithms. Foundations and Trends in Optimization,
1(3):123–231, 2014.

[6] Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. SIAM Journal on Imaging Sciences, 2(1):183–202, 2009.

[7] Amir Beck and Marc Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. IEEE Transactions on Image Processing, 18(11):2419–2434, 2009.

[8] J Webster Stayman, Hao Dang, Yifu Ding, and Jeffrey H Siewerdsen. PIRPLE: a penalized-likelihood framework for incorporation of prior images in CT reconstruction. Physics in Medicine and Biology, 58(21):7563–7582, November 2013.

[9] H Dang, an S Wang, Marc S Sussman, J H Siewerdsen, and J W Stayman. DPIRPLE: a joint estimation framework for deformable registration and penalized-likelihood CT image reconstruction using prior images. Physics in medicine and biology, 59(17):4799–826, September 2014.

[10] David L Donoho, Arian Maleki, and Andrea Montanari. Message Passing Algorithms for Compressed Sensing: II. Analysis and Validation. Phase Transitions, 0:5, 2009.

[11] Jakub Koneˇcny´, Jie Liu, Peter Richtarik, and Martin Tak´aˇc. Mini-Batch Semi-Stochastic Gradient Descent in the Proximal Setting. Pages 1–12, 2015.

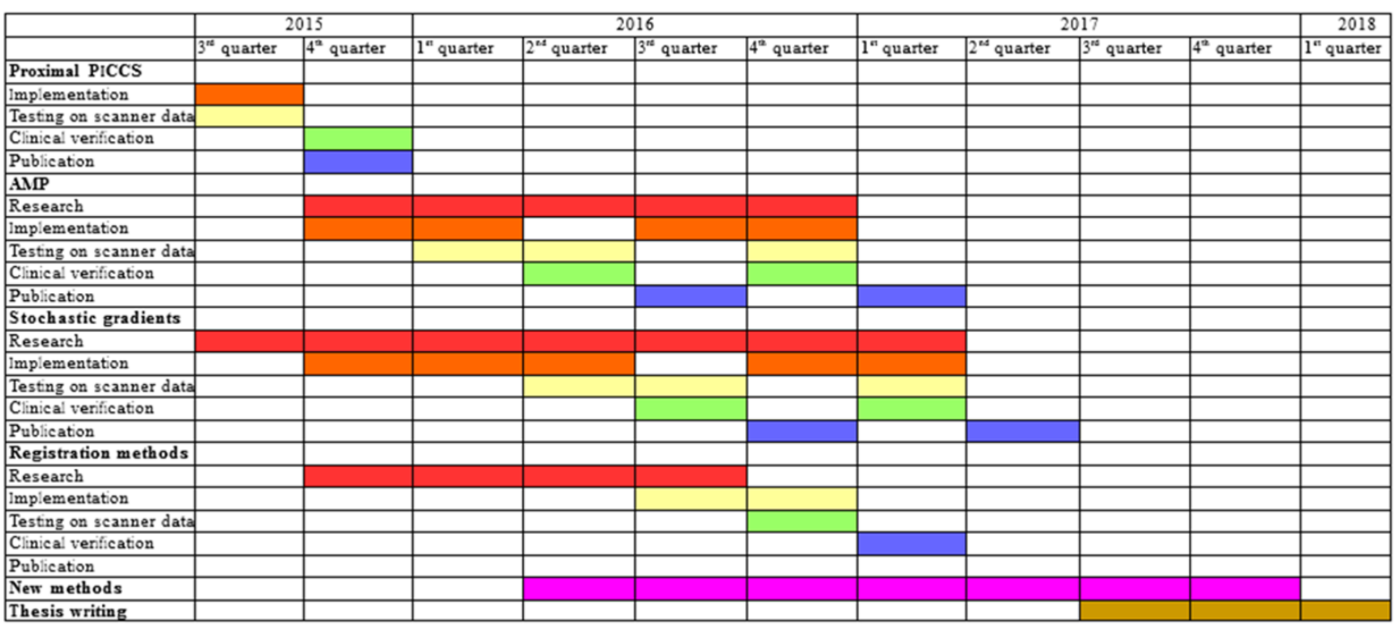

Ashley Hughes
Engineering (Digital Communications) PhD

"The Maxwell Advanced Technology Fund has enabled me to pursue studies in an area which I found interesting as an undergraduate, and wished to continue learning about. Receiving the fund meant that I could study this area full-time, rather than as a secondary field of interest whilst working. Indeed, my having continued studying has directly led to an unsolicited job offer in the field, which I have accepted."

Research Project:
Acoustic Source Localisation and Tracking Using Microphone Arrays

Abstract
A thesis submitted for the degree of Doctor of Philosophy, The University of Edinburgh
June 2015

This thesis considers the domain of acoustic source localisation and tracking in an indoor environment. Acoustic tracking has applications in security; human-computer interaction, and the diarisation of meetings. Source localisation and tracking is typically a computationally expensive task, making it hard to process on-line, especially as the number of speakers to track increases. Much of the literature considers single-source localisation, however a practical system must be able to cope with multiple speakers, possibly active simultaneously, without knowing beforehand how many speakers are present. Techniques are explored for reducing the computational requirements of an acoustic localisation system. Techniques to localise and track multiple active sources are also explored, and developed to be more computationally efficient than the current state of the art algorithms, whilst being able to track more speakers.

The first contribution is the modification of a recent single-speaker source localisation technique, which improves the localisation speed. This is achieved by formalising the implicit assumption by the modified algorithm that speaker height is uniformly distributed on the vertical axis. Estimating height information effectively reduces the search space where speakers have previously been detected, but who may have moved over the ground-plane, and are unlikely to have significantly changed height. This is developed to allow multiple non-simultaneously active sources to be located. This is applicable when the system is given information from a secondary source such as a set of cameras allowing the efficient identification of active speakers rather than just the locations of people in the environment.

The next contribution of the thesis is the application of a particle swarm technique to significantly further decrease the computational cost of localising a single source in an indoor environment, compared the state of the art. Several variants of the particle swarm technique are explored, including novel variants designed specifically for localising acoustic sources. Each method is characterised in terms of its computational complexity as well as the average localisation error. The techniques' responses to acoustic noise is also considered, and they are found to be robust.

A further contribution is made by using multi-optima swarm techniques to localise multiple simultaneously active sources. This makes use of techniques which extend the single-source particle swarm techniques to finding multiple optima of the acoustic objective function. Several techniques are investigated and their performance in terms of localisation accuracy and computational complexity is characterised. Consideration is also given to how these metrics change when an increasing number of active speakers are to be localised.

Finally, the application of the multi-optima localisation methods as an input to a multi-target tracking system is presented. Tracking multiple speakers is a more complex task than tracking single acoustic source, as observations of audio activity must be associated in some way with distinct speakers. The tracker used is known to be a relatively efficient technique, and the nature of the multi-optima output format is modified to allow the application of this technique to the task of speaker tracking.

Chunli Guo
Engineering (Digital Communications) PhD

Research Project:
Compressed Sensing with Approximate Message Passing: Measurement Matrix and Algorithm Design

Abstract
A thesis submitted for the degree of Doctor of Philosophy, The University of Edinburgh
November 2013

Compressed sensing (CS) is an emerging technique that exploits the properties of a sparse or compressible signal to efficiently and faithfully capture it with a sampling rate far below the Nyquist rate. The primary goal of compressed sensing is to achieve the best signal recovery with the least number of samples. To this end, two research directions have been receiving increasing attention: customizing the measurement matrix to the signal of interest and optimizing the reconstruction algorithm. In this thesis, contributions in both directions are made in the Bayesian setting for compressed sensing. The work presented in this thesis focuses on the approximate message passing (AMP) schemes, a new class of recovery algorithm that takes advantage of the statistical properties of the CS problem.

First of all, a complete sample distortion (SD) framework is presented to fundamentally quantify the reconstruction performance for a certain pair of measurement matrix and recovery scheme. In the SD setting, the non-optimality region of the homogeneous Gaussian matrix is identified and the novel zeroing matrix is proposed with an improved performance. With the SD framework, the optimal sample allocation strategy for the block diagonal measurement matrix are derived for the wavelet representation of natural images. Extensive simulations validate the optimality of the proposed measurement matrix design.

Motivated by the zeroing matrix, we extend the seeded matrix design in the CS literature to the novel modulated matrix structure. The major advantage of the modulated matrix over the seeded matrix lies in the simplicity of its state evolution dynamics. Together with the AMP based algorithm, the modulated matrix possesses a 1-D performance prediction system, with which we can optimize the matrix configuration. We then focus on a special modulated matrix form, designated as the two block matrix, which can also be seen as a generalization of the zeroing matrix. The effectiveness of the two block matrix is demonstrated through both sparse and compressible signals. The underlining reason for the improved performance is presented through the analysis of the state evolution dynamics.

The final contribution of the thesis explores improving the reconstruction algorithm. By taking the signal prior into account, the Bayesian optimal AMP (BAMP) algorithm is demonstrated to dramatically improve the reconstruction quality. The key insight for its success is that it utilizes the minimum mean square error (MMSE) estimator for the CS denoising. However, the prerequisite of the prior information makes it often impractical. A novel SURE-AMP algorithm is proposed to address the dilemma. The critical feature of SURE-AMP is that the Stein's unbiased risk estimate (SURE) based parametric least square estimator is used to replace the MMSE estimator. Given the optimization of the SURE estimator only involves the noisy data, it eliminates the need for the signal prior, thus can accommodate more general sparse models.

Development and Alumni
The University of Edinburgh
Charles Stewart House
9-16 Chambers Street
Edinburgh EH1 1HT

Tel: 0131 650 2240

Email: development@ed.ac.uk
www.friends.ed.ac.uk