EPSRC
Engineering and Physical Sciences Research Council

dstl

April 2013 to September 2015

UDRC – Progress Report

Mike Davies, Yvan Petillot, Janet Forbes

Edinburgh Consortium,

Joint Research Institute of Signal Image Processing

Contents

University Defence Research Collaboration

The University Defence Research Collaboration (UDRC) in Signal Processing is delivered in partnership with Dstl, the University of Edinburgh and Heriot-Watt University (Edinburgh Consortium) and Loughborough University, University of Surrey, University of Strathclyde and Cardiff University (LSSC Consortium).

This work is funded by the MOD and EPSRC.

Objectives

The key objectives are to:

- Develop novel signal acquisition and processing techniques to address the needs of the MOD.
- Develop the theory of networked sensor integration to enable future competitive advantage.
- Apply methods in real and simulated data demonstrating effectiveness of the algorithms.

Vision

To develop an ambitious programme of research which will enhance and build upon existing sensor technologies in defence and will provide integrated multi-sensor systems while simultaneously limiting the data overload and maximising data relevance within the network through data acquisition, processing and sensor management.

Director's Introduction

Research focuses on 6 fundamental areas within this field and each research area is led by a senior academic and supported by other academics, research associates and PhD students. The research is divided into the areas below and as the research advances, overlap and synergies are appearing amongst the programmes of work.

Director, Professor Mike Davies

WP 1 Sparse Representation and Compressed Sensing

- WP 1.1 Efficient subNyquist sampling schemes
- WP 1.2 Compressive imaging with sensor constraints
- WP 1.3 Compressed Sensing, beyond imaging

WP 2 Distributed multi-sensor processing

- WP 2.1 Fusion and Registration
- WP2.2 Distributed Decentralised Detection

WP 3 Unified Detection, Localization, and Classification (DLC) in complex environments

- WP 3.1 Estimating targets in scenarios with spatio-temporally correlated clutter
- WP 3.2 Physical Modelling for DLC
- WP 3.3 Man-made object detection

WP 4 Context-driven Behaviour Monitoring & Anomaly Detection

- WP 4.1 Detecting anomalous behaviour in audio-video sensor networks
- WP 4.2 Mobile vehicle monitoring, resource allocation and situational awareness

WP 5 Network enabled sensor management

- WP 5.1 Hierarchical sensor management for target tracking
- WP 5.2 Computationally tractable solutions
- WP 5.3 Multi-objective sensor management

WP 6 Efficient Computation of Complex Signal Processing Algorithms

- WP 6.1 Efficient parallelization of Sensing Processing
- WP 6.2 Implementation of Distributed Signal Processing Algorithms
- WP 6.3 Algorithm/computation resource management

We are now at the half way point of the UDRC research programme and this report gives an update on the progress, highlights and the future focus of the research within the Edinburgh Consortium. This report also informs the mid-term review process of the significant progress against our original grant proposal. The first year of UDRC established an exciting and challenging research programme with the creation of a team with diverse expertise in the field of signal processing. The second year has built steadily on this and developed world leading focused research in signal processing for defence addressing critical gaps for defence application. During this period, the Edinburgh Consortium have produced 69 papers in total, made up of 1 patent, 14 journal papers, 47 conference papers with a further 7 papers submitted or due to be submitted.

This year we have asked for our Strategic Advisory Group to comment more formally on the research. As a result of SAG fulfilling this request, they have delivered an essential role in the provision of guidance and direction on the research programme. The results provided will lend support to our future research strategy for the latter half of the programme timeline. Research is advancing well and the highlights of the work packages can be seen below. Within the highlights, we have listed our most prominent publications, a full detailed list of all publications can be found towards the end of the document.

I hope you enjoy reading our report covering the period from April 2013 to September 2015.

Mike Davies
UDRC Director, September 2015

Research Programme Highlights

WP1 Sparse Representations and Compressed Sensing

A sub-Nyquist Radar ES framework has been developed which in a realistic Electronic Surveillance (ES) simulation setting, outperformed a canonical rapidly swept super heterodyne receiver (RSSR) by up to 20 percent true positive detection for a fixed false positive rate [P1-P4, P11, P15].

In collaboration with WP6, fast GPU implemented Synthetic Aperture Radar (SAR) imaging algorithms based on decimation strategies and parallel processing have been developed providing 330 speed-up over single core standard back projection [P7 and P8].

Proposed autofocus SAR imaging using the sparse nature of dominant targets in scenarios where only partial data is acquired. The proposed algorithm is 20dB better than post-processing with a traditional phase gradient approach [P5].

A new fast spectral decomposition toolbox has been produced for Raman Spectroscopy based on the signal sparsity and non-negativity. It is capable of separating component fractions down to 5% and can also identify the presence of unknown substances from the spectra. Current run time is approximately 1 second on a single core of a 2.0 GHz Intel Xeon processor [P9, P10, P16].

WP2 Distributed Multi-Sensor Processing

We developed the concept of separable likelihoods for distributed estimation of unknown parameters in state space models. Using this framework, we proposed a cooperative sensor self-registration algorithm for distributed fusion networks. Using this algorithm, the platforms can locate themselves in GPS denying environments, based on detections they collect from non-cooperative targets and without transmitting these measurements in the network to other sensor platforms [P21, P19, P25].

We developed an online algorithm for simultaneous localisation and tracking in clusters of sensors (e.g., clusters of bearings only sensors) in GPS denying environments. This algorithm is for centralised processing at the cluster head and features scalability with the number of sensors. It is also capable of accounting for additional location information provided by received signal strength measurements at the receiver front-end of the fusion centre [P22].

In collaboration with Edinburgh Consortium WP5, we developed PHD and CPHD filtering algorithms that are capable of computing the level of confidence for the estimates of number of targets in arbitrarily selected regions. This quantity has the potential to enable regional information-based decision making in sensor network management [P20].

WP3 Unified Detection, Localization and Classification

A complete theoretical framework for sonar Multiple-Input Multiple-Output (MIMO) systems has been developed. By highlighting the theoretical differences between sonar and radar MIMO, new avenues for the application of MIMO systems for recognition, tracking and imaging over a wide range of underwater acoustic problems have been explored [P28, P30, P33].

A new MIMO formulation for broadband MIMO sonar systems which includes a statistical framework for large MIMO sonar has been developed [P33] and have demonstrated that large coherent MIMO systems can provide super-resolution imaging by de-correlating the scatterers present in a single resolution cell. We have also shown that the number of scatterers in a cell can be estimated very accurately and that can lead to automatic target detection and recognition as the majority of objects of interest only have a few important scatterers.

Explored the MIMO design requirement for fully independent MIMO array [P26]. Design parameters and their impact on the correlation of the various views offered by MIMO systems have been studied in detail and compared to equivalent Synthetic Aperture Sonar systems. The MIMO system proposed is a coherent one and requires precise positioning of the transmit and receive elements. Autofocus algorithms for auto-calibration of the MIMO array have been developed [P29]. As a by-product, we have demonstrated that automatic target depth, speed and orientation estimations can be achieved using well-chosen delays between the transmit signals at each element.

The ability to track moving target using defocused time-reversal algorithm has been presented. To validate our theoretical models, a 3D MIMO physics based simulator was developed [P27]. Recently, International trials with DRDC Atlantic and Dstl on novel low and medium frequency systems for Mine Counter Measures (MCM) were organised. In these trails, 3 systems were run together in the same test range were several targets were deployed. The first system was a state of the art SAS system (high frequency), the second system was HWU wideband sonar (medium range frequency) and the last system was the DRDC ex-raise system (very low frequency). A unique dataset spanning an urge frequency band was acquired and will be used shortly to analyse the benefits of each system for target detection and recognition.

Single photon counting LIDAR systems offer the possibility of getting data from targets even with a very low photon count and are ideally suited to detect targets in difficult environments (for instance under foliage) [P36, P37]. We have worked on multispectral LIDAR for the detection of targets casted as an anomaly detection problem. The use of multispectral systems enables the detection of anomalies both in term of discrepancy of the 3D signature but also in terms of the spectral response of the potential target with respect to the surrounding terrain. To date, we are evaluating performance on the detection of several man-made objects and anomalous spectra hidden in a dense clutter of vegetation. As a by-product, our analysis may also allow tree species classification.

WP4 Context Driven Behaviour Monitoring and Anomaly Detection

Novel research on machine learning of target context, looking at social, spatial and temporal aspects, has demonstrated that subtle anomalies enacted by people in surveillance may be detected by the use of advanced visual features, including “gazing” patterns and scene pattern-of life (POL) (as well as more basic features). Our work includes: The development of “intentional priors” which aggregate spatial and temporal priors over multiple features. This has unified the “contextual” and “pattern-of-life” into basic target tracking [P39, P41].

A technique for detecting social grouping in crowded surveillance scenes [P42] and a novel algorithm for on-line anomaly detection and learning of target pattern-of-life from Wide Area Motion Imagery (WAMI) and other ‘big data’ sources has been developed.

A new head-pose classification algorithm based on Deep Learning that significantly out performs the prior state-of-the-art [P50, P51].Research on joint audio-video localisation and tracking has shown that benefits can be achieved even with low complexity systems. Work includes the development of a multi-modal tracker capable of tracking through occlusion [P40, P46, P47], an algorithm for dominant sound source localisation using movement [P43] and a new theory for performing multi-person activity recognition and anomaly detection without requiring large training corpora [P38]

WP5 Network Enabled Sensor Management

Development of unified estimation framework for detection, tracking, and classification of a unknown number of objects in cluttered environments from heterogeneous sensors, design and implementation of a multi-object filter derived from this framework [P52]. Application on a maritime surveillance scenario involving target classification[P27].

Incorporation of sensor management policies to multi-object estimation framework [P20]. One solution based on population activity for tracking algorithms derived from Finite Set Statistics framework or novel unified estimation framework [P62]. One solution based on information-theoretic decisions for tracking algorithms derived from novel unified framework, but adaptable to traditional track-based approaches [P53].

Exploitation of the novel estimation framework for a space situational awareness scenario with a Doppler radar, in view of the integration of the UK ground-based assets for space surveillance (radars, telesopes, lidars, etc.) in a single estimation algorithm [P60]. Exploitation of the estimation framework for sensor calibration in multiple-target detection and tracking scenarios [P59, P61].

WP6 Efficient Computation of Complex Signal Processing Algorithms

The accuracy and reliability of state-of-the-art classifiers on Synthetic Aperture Radar (SAR), Synthetic Aperture Sonar (SAS) and visual datasets have been evaluated. We detect and classify a variety of objects including pedestrians and mine-like shapes. We have investigated methods for improving reliability, and can demonstrate the ability to detect objects belonging to previously unseen classes; this is important for tasks requiring human operators and high-level decision algorithms which will act based on the results from object detectors [P64, P66].

A pedestrian detector which responds to uncertain images with less confident prediction scores has been produced. This is an improvement over existing state-of-the-art detectors, which are often extremely overconfident in the presence of uncertain samples. A Graphics Processing Unit (GPU) implementation of a Gaussian Processes Classifier was written and performed favourably when compared to existing approaches [P64, P68].GPU-accelerated implementations of state-of-the-art image formation algorithms for SAR imaging have also been developed [P8].

Computer vision algorithms for distributed anomaly detection and object recognition have been implemented on an Android smartphone. A client-server architecture has been used to investigate uploading images or features from a smartphone to a server for faster processing, and when it is most power-and time-efficient to do so [P69].

A simulator for evaluating distributed localization algorithms has been developed. This allows us to implement and test methods for detecting and localizing hostile Radio Frequency (RF) emitters in scenarios with limited processing power and communication bandwidth available. We are using this to develop and improve localization algorithms used on devices running in a mobile ad-hoc network (MANET), with the objective of minimising localization error, power consumption and the volume of data transmitted over a ‘friendly’ distributed network.

Management

The UDRC has established an effective management structure with formal processes to manage issues, research collaborations and communication. The Edinburgh Consortium is managed by Director, Mike Davies and Deputy Director Yvan Petillot. They meet on a monthly basis along with the Project Manager Janet Forbes to review progress and outputs.

Internal meetings are held quarterly and allow the academics, researchers and the project management team to update on progress, discuss the way forward so ensuring the efficient and effective implementation of the UDRC research. The researchers meet fortnightly to discuss the status of work, collaboration opportunities and future perspectives. These meetings serve as a forum for the researchers to communicate their work and ideas in order to foster collaboration between different work packages. UDRC researchers have also established a journal reading club which will catalyse the interactions further.

In May the researchers and the academics met for an away day activity where each group will answer a set of predetermined questions about the research and its desired impact. The outcomes from this activity has fed into the mid-term review process and the future research strategy for the remaining 2.5 years.

A data repository has been set up to allow for the storage of data, research and software allowing both universities to work together more efficiently. This area allows for the storage of large amounts of data. A wiki area has also been established. This ensures that all important documents relating to UDRC are easily accessible. Documents in this area include presentations, minutes, instructions and templates.

Working closely with Dstl is an essential part of the project. Each of the work packages have been given Dstl exploitation points of contact and meetings are set up regularly to allow a two way exchange of information to ensure that research heads in the right direction and potential opportunities are exploited.

An enabling contract between Dstl, University of Edinburgh and Heriot-Watt University has been drafted, this has led a number of contracts with Dstl which has taken work further including projects on Raman spectrometry and underwater and space tracking.

In order for Dstl to review outputs and progress, both the Edinburgh and LSSC Consortium Directors meet at Dstl for a 6 monthly Governance meetings to review the project.

The Strategic Advisory Group (SAG)

The Strategic Advisory Group (SAG) meet every 6 months; the purpose of this group is to provide feedback to the consortium on the research it is performing and to ensure that the research output continues to target areas of national and international importance for the defence sector.

The value of the research is judged with regard to MoD priorities, academic world class quality, relationship with other MoD and UK industrial research and the relevance to the needs of the UK defence industry (including practicality of implementation). The SAG also provides feedback on the engagement activities, judging them on the facilitation of stronger links between signal processing research groups, defence industries and the provision to the government defence sector. The members also advise on the development of a single Community of Practise for defence-related signal processing research, spanning academia, industry and government.

This year we have asked for our Strategic Advisory Group to comment more formally on the research. These comments are very important and the results from this review coupled with the away day discussion will lead to the development of our strategy for the remaining programme of research.

SAG Membership

Angus Johnson, Thales UK

Andrew Baird, Dstl

Alfred Hero, University of Michigan

Jonathan Evans, SeeByte Ltd

Dean Thomas, Roke Manor Ltd

Jordi Barr, Dstl

Stephen Clark, Selex ES

Malcolm MacLeod, QinetiQ

Nigel Birch, EPSRC

Paul Thomas, Dstl

Henry White, BAE Systems

Ros Knowles, Dstl

Research Team

The Edinburgh Consortium comprises signal processing experts from the University of Edinburgh and Heriot-Watt University and is one of the two Consortiums funded for phase 2 of the University Defence Research Collaboration (UDRC).

The Joint Research Institute of Signal and Image Processing is a partnership between these two universities and incorporates the activities of:

- Institute for Digital Communications (IDCOM), University of Edinburgh.
- The Research Institute of Signals, Sensors and Systems (ISSS), Heriot-Watt University.

Academic Staff

Prof Mike Davies
Prof Yvan Petillot
Prof Bernie Mulgrew
Prof John Thompson
Prof Andy Wallace
Dr Neil Robertson
Dr Daniel Clark
Dr James Hopgood
Dr Mathini Sellathurai

Research Associates

Dr Mehrdad Yaghoobi
Dr Murat Uney
Dr Yan Pailhas
Dr Rolf Baxter
Dr Emmanuel Delande
Dr Calum Blair
Dr Jeremie Houssineau
Dr Salvatore Caporale

Research Students

Di Wu
Puneet Chhabra
Jose Franco
Saurav Sthapit
Kimim Kim
Alessandro Borgia
Alexey Narykov

Project Management

Janet Forbes
Madeleine McBeath
Audrey Tainsh

WP1 Sparse Representations and Compressed Sensing

Research Leader: Mike Davies

Academics: Bernard Mulgrew, Mathini Sellathurai

Research Associate: Mehrdad Yaghoobi

PhD Student: Di Wu

The aim of this work package is to explore the potential use of sparse structures in the state-of-the-art signal processing applied to battlefield sensing. While the sparse and compressible signals exist in a number of defence applications, the exploitation of sparsity has not always been realized and is worthy of further investigation. We expect to deliver efficient approaches for practical sensing and imaging scenarios in the specific fields of the Radar Electronic Surveillance Measures (ESM), SAR imaging systems and chemical detection.

In WP1.1, the aim is to use the compressive sampling for the analog to information conversion. The objective is to develop computationally low-cost and robust techniques for ultra-wide band Radio Frequency (RF) signal conversion.

Compressed sensing is used to present a more efficient way for radar imaging. WP 1.2 explores the sensor constraints, including phase ambiguity, calibration, RF interference, using Synthetic Aperture Radar signal structures. As imaging and sensing in defence often deals with a large amount of data, suitable techniques for compressed sensing and sparse representations which can handle these problems, will also be investigated in WP1.2.

Outcomes

There has been progress in sub-Nyquist sampling for radar ES, SAR applications and Raman Spectroscopy. Nine papers were presented in this period.

Our novel compressed sensing Radar ES has been presented in [P2, P3, P4, P11] and the patent application [P1] presents our IP for this technology.

The compressive autofocus work was published in the IEEE transactions [P5] and the combined compressed sensing strategy for Low Frequency SAR has been presented in [P6], as well as in an Invited Lecture at the EuRAD Workshop on Compressed Sensing Radar Applications [P12].

In conjunction with WP6 and motivated by the decimation-in-time FFT algorithms, in [P7] we proposed a novel fast back projection algorithm for SAR imaging with spatially controllable errors compared with other fast back projection algorithms. The use of multi-core processing and graphic processing units to further speed up the proposed algorithm was investigated in [P8].

The initial Raman Spectroscopy work was presented in [P9] to help qualitatively analyse the chemical components and quantitatively estimate the concentrations. One of the fast nonnegative greedy algorithms developed was subsequently reported in the IEEE Signal Processing Letter [P10].

In related work Mathini Sellathurai has worked on the Electronic Surveillance of communications signals yielding two papers [P13, P14]. In [P13] this work considered developing real-time programmable low complexity techniques for blindly detect various OFDM modulated passive communications signals (recognize and estimate parameters of OFDM signals) and in complex channel environments and low (negative) SNR conditions. The developed techniques were tested using HWU's NI wireless testbed in the real world environment. [P14] proposes a novel way of extracting communication signals from a noisy spectrogram using OFDM computer vision and neural network based techniques.

Progress

We investigated the sub-Nyquist sampling for radar ES, sparsity-driven techniques for Raman Spectroscopy and various tasks in Synthetic Aperture Radar applications.

In Y1, we identified an efficient approach for ESM based on efficient sampling. Three approaches, namely Random Demodulator (RD) [R1], Modulated Wideband Converter (MWC) [R2] and Multicoset Sampling (MS) method [R3], were identified as the most successful settings. Our focus was to present an efficient DSP unit for this application. Also we achieved a good progress in the sparsity based SAR imaging, and started work on Raman spectral deconvolution.

We began by choosing the analogue hardware proposed in multicoset sampling. This provides a bank of delayed and sampled signals at an overall rate much lower than Nyquist (the coset signals). The proposed framework then relies on digital fractional delays (DFD)'s, followed by a time-frequency (TF) transform, and a simple subband classifier. The whole process is non-iterative and therefore could be implemented in a real time pipeline.

We subsequently investigated the optimal choice of channel numbers and delays; these are chosen to form a so-called Grassmannian frame, in order to maximise the detection probability and output SNR. The second major advance was the inclusion of the DFDs within a modified TF transform. Apart from computational savings this also achieved a boost in the output SNR of between 0.5 to 1 dB.

In our Radar research we continued our work on compressed sensing for SAR imaging, following on from the UDRC phase 1 project. A key development here was the incorporation of autofocus techniques within the iterative reconstruction algorithm base around dictionary learning techniques. The additional cost was negligible as the phase correction term could be implemented at each iteration in closed form.

At the end of Y1 a new PhD student, Di Wu, started and initially begun investigating sparsity based super-resolution, with particular application to joint SAR+ GMTI imaging.

In Y2, we considered comparisons between our proposed subNyquist Radar ES system and made popular wideband electronic surveillance technique in the industry, such as Rapidly Swept Superheterodyne Receivers (RSSR). In SAR, we continued our work on multi-channel SAR for SAR + GMTI processing and have begun to explore the practicalities of exploiting sparsity in 3D Low Frequency SAR imaging. We also began a separate project on signal decomposition for Raman spectral analysis based on a Dstl challenge presented at one of the UDRC themed meetings.

In compressed sensing Radar ES, we analysed our LoCoMC algorithm in a CFAR framework and derived the optimal thresholding parameters in the radar Electronic Surveillance (ES) which helped optimise the processing gain of the receiver. We also extended the sub-Nyquist Radar ES simulations using a larger set of pulses to provide a good statistical analysis.

For the SAR work we have received the AFRL Gotcha GMTI data set and started to look at the realistic SAR/GMTI problem. The challenge is to realise simultaneously the target detection, velocity estimation and SAR imaging in the presence of urban clutter. We have begun to frame the problem as an undersampled missing data problem and incorporate multi-channel balancing techniques [R4] and moving targets relocation effects [R5]. These factors are among the most important principles in SAR/GMTI, yet the optimal methodology has not been established in the SAR community.

In the Dstl project, Raman spectral deconvolution was investigated to analyse and separate the components of chemical mixtures, and quantitatively determine the concentrations by utilising the sparsity of components, compared to the whole chemical library. For this work Dstl have provided further data to evaluate our algorithms in more challenging scenarios. For the processing we have developed two exceptionally fast non-negative greedy sparse approximation techniques to accelerate the deconvolution and separation process and provided Dstl with a Matlab toolbox for this work. In future work it is anticipated that we will work with Dstl and their industrial partners to help refine the system and port the algorithms onto the sensors on board ARM processor.

Future Direction

In the forthcoming year we plan to focus on the advanced SAR imaging projects. There will be two components. The first will investigate sparsity driven SAR + GMTI algorithms [P17], while the second will look at the challenging problem of 3D SAR image formation from a multiple passes. Dstl and industry have identified volumetric SAR imaging as a key research priority, extending beyond conventional approaches, e.g. interferometric and tomographic, and solving the true 3D reconstruction of the scene, using advanced fast reconstruction techniques from limited trajectories.

The issue of out-of-focus data becomes more challenging, when the recorded pulses have to be combined for 3D SAR. We have already started to investigate the problem of phase error correction, as well as other issues in SAR imaging, for example defocusing, moving objects and non-calibrated equipment, which can be interpreted as a phase correction problem. This will be particularly interesting for LF imaging where the capability to view through foliage and into buildings is of interest. We have received the Bright Sapphire LF SAR trials data from Dstl and are currently in discussions with QinetiQ on its use for LF 3D imaging. We also plan to prioritise video SAR processing (PhD student), including the development of algorithms for image sequence generation and the incorporation of dynamic imaging such as GMTI within the image sequence.

The dynamic, LF and 3D SAR research are built on our previous compressed sensing SAR work and are making heavy use of the fast algorithms developed in conjunction with WP6.

We will also continue our work on sparse spectral decomposition. The next phase will involve porting the algorithm onto a portable low-power device. We also plan to investigate the possibility of incorporating task specific extensions, e.g. preferential spectral decomposition. We would also like to extend the idea to more general hyperspectral image decomposition.

WP2 Distributed Multi-Sensor Processing

Research Leader: Bernard Mulgrew

Academics: Daniel Clark, John Thompson, Neil Robertson, Mathini Sellathurai

Research Associate: Murat Uney, Salvatore Caporale

PhD Student: Kimin Kim

This work package addresses the challenges in detecting and tracking targets with networked sensors of various modalities. In order to meet with the requirements of performance, flexibility and fault tolerance under resource constraints such as limited communication bandwidth and energy, we investigate distributed solutions which avoid a single designated processing centre. We also address challenges in providing scalable solutions in centralised settings to facilitate multi-sensor exploitation.

In WP2.1, Fusion and registration, we developed methods for sensor registration and fusion in networks of sensors. In particular, we considered the problem of finding respective sensor locations and orientations in the absence of direct measurements of these quantities. Instead, we exploited detections from the targets in the surveillance region together with the multi-object estimates shared in the network for distributed fusion.

The second stage of this research, WP2.2 Distributed/Decentralised detection which started in January 2015, considers detection of targets with networked sensors. The algorithms developed in the first part of this work package (WP2.1) make use of the outputs from detection algorithms run locally at the sensor platforms. The aim of WP2.2 is to address challenges in detecting targets by exploiting the diversity and/or the extended coverage provided by having more than one sensor. Detection of manoeuvring and/or dim targets is particularly challenging. We are exploring track before detect strategy which updates target trajectory estimates using the signals output by the receiver front-end processing while accommodating the information from target trajectory/location in the processing chain.

Outcomes

A novel distributed algorithm for sensor self-calibration has been developed. This algorithm is built upon the concept of node-wise separable parameter likelihoods developed in WP2.1 [P21, P25] and demonstrated for sensor self-localisation in fusion networks in GPS denying environments using simulations. A journal article describing the details of the algorithm together with theoretical results on the quality of approximations involved is now accepted for publication in IEEE Transactions on Signal Processing [P19].

A centralised algorithm for simultaneous sensor localisation and target tracking that features scalability with the number of sensors has been created [P22]. This algorithm is useful for tracking targets with clusters of sensors which individually have limited target observability but collectively can provide sufficient information for tracking such as clusters of bearing only sensors.

The concept of regional variance in multiple target tracking [P20] was presented in collaboration with WP5. We developed PHD and CPHD filters capable of computing a level of confidence for the number of target estimates in arbitrary regions.

We have developed a maximum likelihood framework for estimation of the parameters that characterise the matched filter outputs in active sensing which will in turn be input to a detection algorithm [P24]. This work is a building block towards addressing detection in networked radars and similar detection networks. We presented our work at the IEEE Workshop on Statistical Signal Processing 2014, Bayes Lectures 2014, Sensor Signal Processing for Defence Conference 2014, 2015, and in other UDRC events.

Progress

We commenced an investigation into a distributed fusion network scenario motivated by previous work [P18]. In this architecture, sensor platforms perform local filtering of their target detections and exchange the multi-object posterior distributions produced with their neighbours in order to improve upon the local tracking performances. Tracking algorithms used for local processing can estimate target positions as well as the number of targets in a given region. Measures of confidence on the latter, for example, the variance of the target number estimates is not available from any of these algorithms. This information has the potential to enable regional information-based decision making in fusion networks. Estimates of the number of targets in arbitrary regions are typically found via the first order moment of the corresponding population density. This moment function is directly computed by the PHD and the CPHD filters. In collaboration with WP5, we developed versions of these filters which are capable of computing the second order moment of the population density leading to the variance of the target number estimates. We demonstrated this novel concept of regional variance in several examples [P20].

From June 2013 to December 2014, we focussed on distributed sensor registration in fusion networks. In multi-sensor fusion, sensor registration/calibration refers to specifying the parameters necessary to relate the reference frames of different sensors. These parameters include sensor locations and orientations and their full knowledge is needed in order to integrate the information from different sensors [R6]. In some cases, registration parameters can be measured using on board devices such as gyroscopes and/or global navigation space systems (GNSS). In a range of applications, however, such measurements are not available. For example, GNSS cannot be used for underwater fusion networks due to the signal propagation constraints and otherwise are also vulnerable to jamming.

We considered a fusion network in which the sensor nodes perform local filtering of the measurements from the targets in their field of view and communicate the output posteriors to their neighbours in the underlying communication network. The centralised solution for estimating the registration parameters, for example, the sensor locations, involves a parameter likelihood which requires all the sensor measurements collected across the network and over time to be available at the fusion centre [R7].

There are two main challenges in developing a distributed solution in this scenario: The first is the specification of likelihood functions based on the target posteriors transmitted between sensors as opposed to the complete network-wide measurement history in the centralised solution. The second challenge is the identification of a distributed inference method for the estimation of network-wide parameters using these likelihoods. In order to overcome the first difficulty, we developed node-wise separable likelihoods. These likelihoods consider sensor pairs and can be computed locally using local target detections and the incoming posteriors (Figure 1). For the second part of the problem, we adopted a pairwise Markov Random Field model with the developed likelihood functions as the edge potentials and used Belief Propagation for decentralised self-localisation.

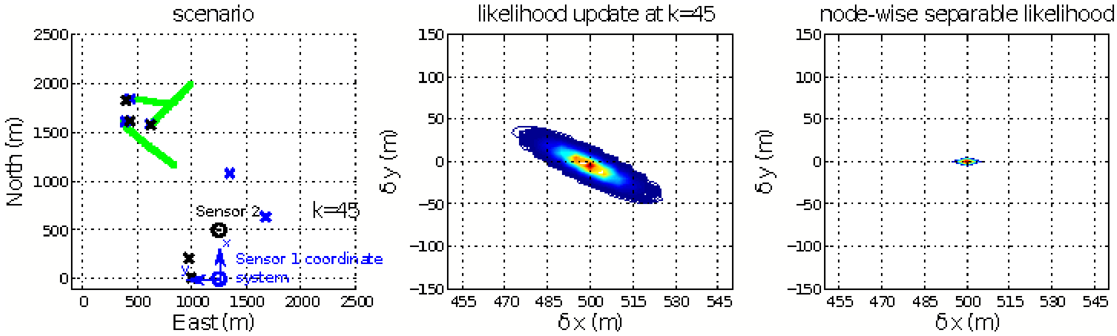

Figure 1. Node-wise separable localisation likelihood for a pair of sensors: (left) 2 range-bearing sensors collect noisy measurements from multiple targets (green tracks) with false alarms and imperfect detection probabilities (black and blue crosses). They perform local multi-target filtering and exchange their posteriors (here, Sequential Monte Carlo (SMC) Probability Hypothesis Density (PHD) filtering [R7] is used and the multi-target posteriors are multi-object Poisson distributions). Starting from time step k=1, we observe the change of the parameter likelihood. (Central graph) The update term at k=45 for the likelihood of the respective location of Sensor 2 in Sensor 1's coordinate system. (right) The node-wise separable likelihood obtained through the updates from k=1 to k=45

The preliminary results [P21] of this approach featured sensor localisation using non-parametric Belief Propagation [R8] with node-wise separable likelihoods based on local Sequential Monte Carlo (SMC) Probability Hypothesis Density (PHD) [R9] filtering of cluttered and noisy multi-target detections with imperfect detection rates.

We developed belief propagation libraries for Gauss Markov Random Fields (MRFs) and non-parametric pairwise MRFs employing particle representations.

We considered a hierarchical architecture for fusion networks (Figure 2): In the first tier, sensor platforms that do not have a sufficient degree of target observability form clusters with cluster heads acting as processing centres for joint filtering of their measurements. In the second tier, decentralised in-network processing takes place among cluster heads and/or sensors with sufficient degree of observability. Our research on distributed registration had considered the second tier problem in 2013. We also considered registration in the first tier, when the target detections across the cluster is to be filtered centrally at the cluster head (or, fusion centre). We developed a centralised simultaneous localisation and tracking algorithm for clusters of bearing-only sensors which feature scalability with the number of sensors. In the case of bearing-only modality, the node-wise separable likelihoods are not guaranteed to provide an approximation to the centralised likelihood with sufficient quality. Instead, we use thinned junction trees in order to achieve scalability with the number of sensors [P22].

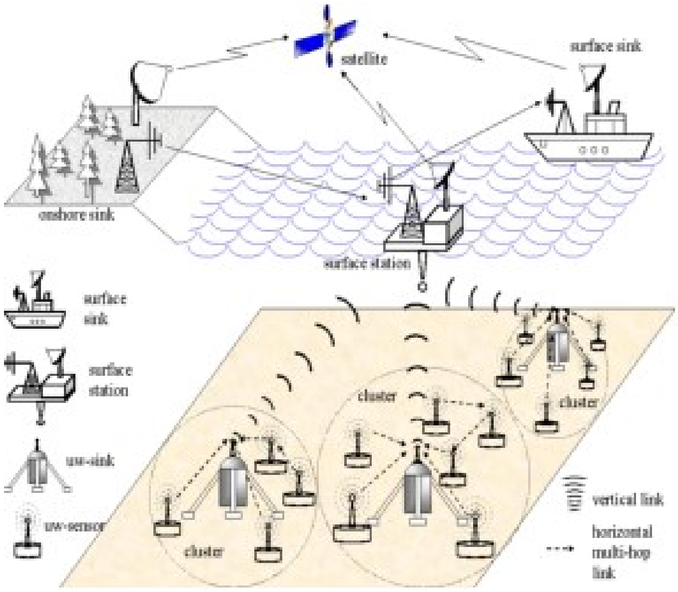

Figure 2. Illustration of an underwater fusion network (from I.F. Akyildiz, D. Pompili and T. Melodia, “Underwater acoustic sensor networks: research challenges,” Elsevier Ad Hoc Networks, 2005

The developed algorithm partitions the problem into sub-problems of a solvable size and then merges the solutions in a coherent fashion with scalable computational complexity. We achieve this by assuming a Junction Tree model for the parameter posterior (Figure 3). This model decomposes the problem and specifies how the partial solutions can be merged (i.e., the Junction tree algorithm) [R10]. From a computational perspective, our algorithm works as a Gibbs sampler for individual sensor locations. Specifically, we window the measurement histories and generate samples from single sensor distributions for consecutive windows. The complexity of the sampler is controlled by selecting the width of the Junction Tree (i.e., the number of variables in the variable nodes). This framework also allows us to exploit additional localisation information such as the received signal strength at the fusion centre for improving the robustness and accuracy of the algorithm (Figure 4-5).

Figure 3. An example Junction Tree model for the locations of bearing-only sensors with respect to the cluster head. (right) A sensor cluster. (middle) An example triangulated Markov graph. (right) The Junction Tree corresponding to the selected triangulation.

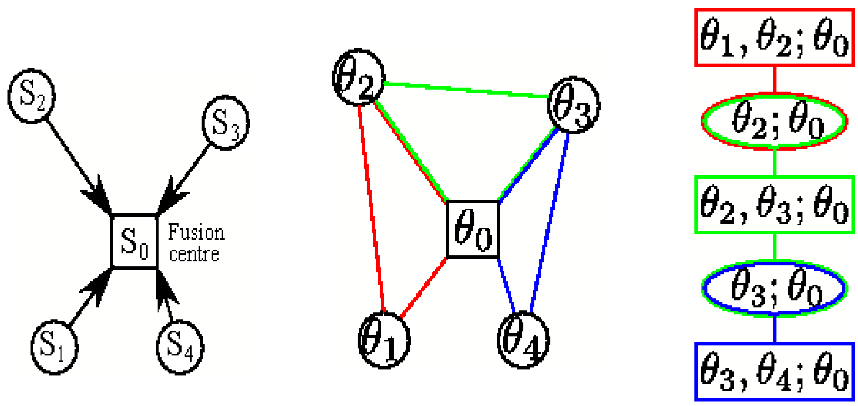

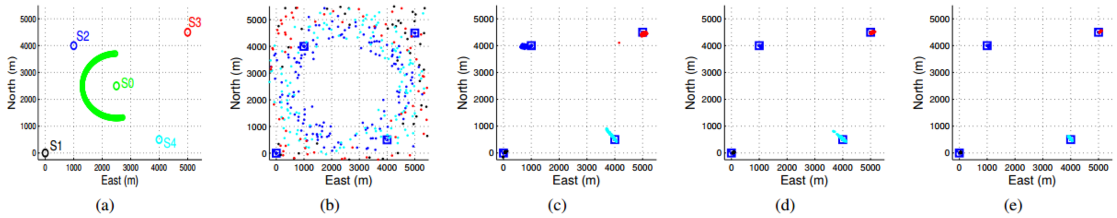

Figure 4. A demonstration of the proposed algorithm: (a) A target (green track) inducing bearing-only measurements on the peripheral sensors (S1-S4) as well as the cluster head (S0). (b) Scatter plot of the particles generated from the location distributions for time window n=1, (c) n=5, (d) n=10, (e) n=15.

Preliminary results of this work was submitted to Sensor Signal Processing for Defence Conference 2014 [P22].

We also considered whether the localisation problem could be solved with sufficient accuracy using receiver front-end statistics at the cluster head such as the received signal strength (RSS) and angle-of-arrival (AoA).

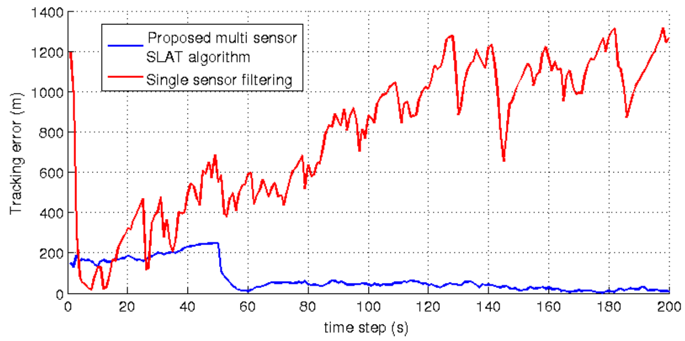

Figure 5. Comparison of the tracking errors for using only the bearing measurements at the cluster head (red) and the proposed online algorithm (black).Note that the particle filter “looses” the track when only the lines of bearing from a single sensor (S0 in Figure 4(a)) is used. Our algorithm makes it possible to use the bearing measurements from other sensors by jointly estimating their locations.

Cramer-Rao Lower Bounds (CRLBs) [R11] was used for these measurements and for the various scenarios, the CRLBs indicate a wide region of uncertainty for sensor locations. Therefore, RSS and AoA fail to lead to a sufficiently accurate tracking performance motivating the use of our target aided approach. One side benefit of this investigation is the provision of further insight to the MANET detection and localisation Dstl challenge.

We also investigated further theoretical results regarding the node-wise separable parameter likelihoods and derived information theoretic error bounds for the node-wise separable likelihoods revealing their quality of approximation to the centralised likelihood for a sensor pair. In particular, we developed the Kullback-Leibler divergence (KLD) between the centralised and the proposed likelihoods in terms of the average Mutual Information (MI) between current measurements and the sensor histories summed with a non-positive term. After expanding the MI terms as differences of the Shannon Entropies of target distributions conditioned on various combinations of the measurement histories, we related the KLD to the filtering performances. For the linear Gaussian case, we have been able to evaluate these expressions exactly, without the need for any Monte Carlo approximations.

An extensive experimentation with our distributed self-localisation algorithm was carried out, in order to develop an understanding of the behaviour of the algorithm for varying clutter rates, probability of false alarms, time window lengths and Brownian motion step size. We have discovered that the algorithm can be made robust to false alarms by an appropriate selection of the Brownian motion step size which acts as a smoothing parameter for the estimated calibration marginals (Figure 6).

Figure 6. Demonstration of the distributed registration algorithm with recursive node-wise separable likelihood updates and BP messaging: The nodes of the network (upper left) perform local filtering (SMC-PHD [R9]) of the target measurements (coloured crosses are detections from the targets with green tracks) and exchange multi-target posterior with their neighbours. For a selected time window T, they update node-wise separable localisation likelihoods. These local updates are in linear complexity with the number of detections. At the end of T, they iterate BP messaging for a selected number of steps (here, we use Non-parametric Belief Propagation [R8] which is a particle algorithm compatible with the filtering performed). Here, the scatter plots of the particles generated from the local position densities are presented over time (upper left through lower right).

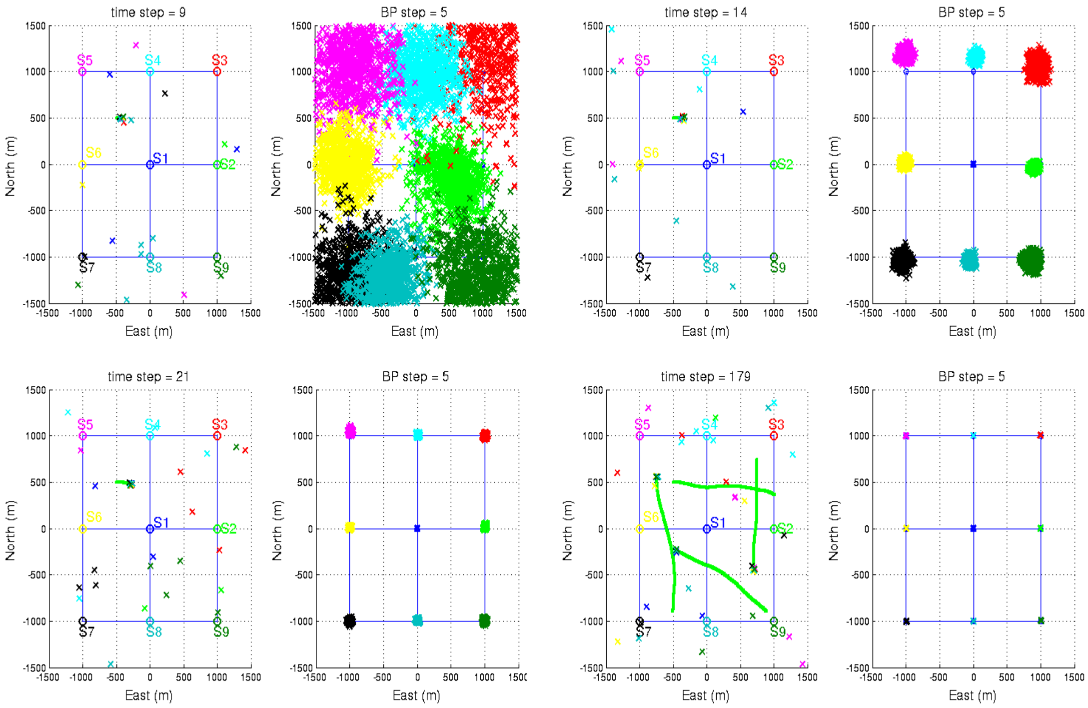

Kimin Kim started on his Ph.D. research project in the context of WP2.2. The goal of the project is to address challenges in detecting targets in wide area surveillance networks. A motivating example of such systems is widely separated multiple-input multiple-output radar. This research has started exploring adaptive processing techniques for long time coherent integration and will explore track before detect strategies in monostatic and multi-static sensor systems. A parallel line of investigation has been commenced on track before detect in the presence of multiple targets and unknown noise characteristics with a particular interest in multiple input single output sensor settings.

Future Direction

The future work in WP2 covers sub-work package WP2.2 distributed detection. We will have two parallel lines of research in distributed detection. One line of investigation will explore space time adaptive processing strategies for single receiver arrays with single and multiple transmitter settings. This research is pertinent to detection in single radars, as well as MIMO systems from a collection of multiple input single output (MISO) systems point of view.

A complementary line of research will investigate methods for efficiently communicating and combining likelihoods for “target plus interference” and “interference alone” within the network. This is a key element in characterising the trade-offs in distributed/decentralised detection and addressing this problem in networks of radars.

WP2.1 remains open for possibilities of demonstrating the calibration algorithms developed on real/simulated data from Dstl and/or other programme stake holders and for completing a journal article covering SSPD 2014 work on scalable simultaneous localisation and tracking for sensor clusters [22].

WP3 Unified Detection, Localization and Classification

Research Leader: Yvan Petillot

Academics: Daniel Clark, James Hopgood, Andrew Wallace

Research Associate: Yan Pailhas, Salvatore Caporale

PhD Student: Puneet Chhabra, Jose Franco

This aim is to understand and model difficult and complex environments. Traditional algorithms for detection, classification or identification are based on simplistic models of noise, clutter or multipath. Consequently, most of them fail to achieve useful or meaningful results in complex maritime environments. We aim to develop realistic physical based models for the full sensing chain from the sensors themselves to the complex interaction with clutter/target and the propagation in the environment. A physical understanding of the clutter rather than ad hoc and simple statistical models will help to develop new DLC (Detection, Localisation and Classification) algorithms with optimal performances and reduced computational power as well as in situ environment adaptability for greater robustness.

Outcomes

The foundation for MIMO sonar systems has been investigated and developed looking at the feasibility and great potential for many applications including harbour surveillance. MIMO systems have great flexibility and we have been able to address specific applications such as speed and depth estimation as well as automatic acoustical target tracking.

Over the first half of the project, a total of 13 publications have been submitted in several international conferences and journals. The wide dissemination had brought to MIMO sonar a better visibility and understanding and also work related to the proposal of several autofocus algorithms, which allow tracking of a target, estimation of its depth and/or its speed and heading automatically.

In WP3.3 we have developed an anomaly detection algorithm that considers Multi-Spectral (MS) FW Ladar (MSL) measurements as a set of multi-dimensional data samples [P36]. The framework allows the detection of spectral and temporal anomalies in FW-MSL data. In the signal domain we define an anomaly as a full waveform time and spectral signature that does not conform to a prior expectation, defined using a learnt subspace (dictionary) and co-occurring local-patterns.

We have developed an optimization algorithm for subspace learning based on stochastic approximations and augmented our objective function with a discriminative term that represents the subspace's separability properties. This algorithm is now being tested on the simulated data.

Progress

Our main focus is on underwater acoustic sensors. MIMO sonar systems are the natural extension to more classical sonar systems. MIMO also offers straightforward solutions to numerous practical challenges [P26, P28, P29]. One example of this is the Dstl challenge 38, Reliable automated detection and identification of underwater objects using unmanned sensor systems. For this challenge, a large MIMO sonar systems was studied with an emphasis on harbour protection.

So far MIMO has exclusively been studied for radar systems. Several models and formulations have been proposed for radar, which stipulate the multi-static relation between transmitters and receivers through the channel matrix. Traditionally the channel matrix models both wave propagation and target scattering. Underwater acoustic wave propagation however is considerably more complex than atmospheric Electromagnetic (EM) wave propagation especially in shallow water environment. For this reason, we produced a general MIMO sonar model suitable for narrow and wideband systems where propagation and scattering are clearly differentiated [P33]. We can now describe a MIMO sonar system by the following equation:

$$z_{lk} = H_{lk}(X_0, w)F_{\square}(w, \theta_l, \phi_k)S_k(w)$$

If radar and sonar systems are both coherent systems, there are fundamental differences, especially in the scattering process. Man-made object sonar echoes, in particular, can be accurately modelled by a finite and relatively small number of scatterers. By using this property and assuming a large and independent MIMO array, we demonstrated that large MIMO sonar systems have "*built-in*" recognition capabilities [P26]. We demonstrated, for example, that it is possible to assess from the target MIMO response the number of scatterers of a target of interest (Figure 7). The second fundamental result of statistical MIMO system is the super resolution capability [P28]. Thanks to the following convergence:

$$\lim_{N\to\infty} N\Gamma(Nx, N, 1) = \delta(1-x)$$

We demonstrated that large MIMO systems can solve the speckle and then separate extremely close scatterers (Figure 8). We also established that, from an imagery point of view, coherent MIMO sonar system surpasses the state of the art SAS imagery [P29].

We showed the importance of independent views for MIMO systems and proposed a formal definition of independent views. We then stated the MIMO intra-views distance correlation matrix derived from the distance correlation metric from Szekely [R12]. This matrix can be seen as a figure of merit for MIMO sonar system design (Figure 9). Maximising this figure of merit during the MIMO design process is essential to unlock the statistical MIMO capabilities.

We then developed a 3D MIMO simulator capable of generating MIMO time signals from a full 3D environment. The simulator includes multiple transmitters and receivers separated spatially, bi-static reverberation contributions (bi-static reverberation levels are computed using bi-static scattering models developed by the Applied Physics Laboratory, University of Washington (APL-UW)), multi-paths (bottom and surface bounces), fractal seabed elevation map; different types of seabed sediments, mid-water targets (Figure 10).

The 3D MIMO simulator enabled several harbour protection scenarios to be built including autonomous underwater vehicle (AUV) intrusion in busy and cluttered environment [P30]. We worked closely with Jeremie Houssineau and demonstrated the capability for MIMO sonar systems to successfully identify and track mid-water targets in complex environment using the Hypothesised multi-object filter for Independent Stochastic Population (HISP) filter recently developed (Figure 11).

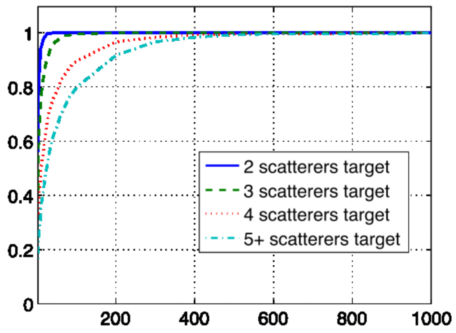

Fig. 7: Correct classification probability against the number of independent views for 4 classes of targets (2, 3, 4 and 5+ scattering points targets).

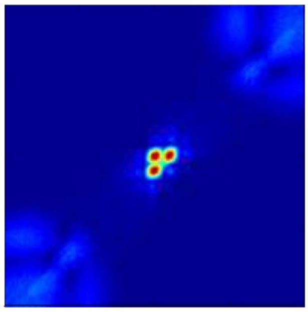

Fig. 8: 3 scatterers target MIMO image using 10 transmitters and 10 receivers with 3 metres spacing. The scatterers are one λ apart.

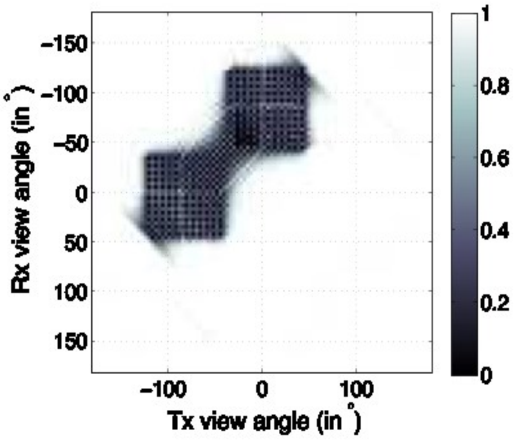

Fig. 9: Full MIMO inter-correlation distance matrix for a 10 Tx and 10 Rx MIMO array with 3 m element separation.

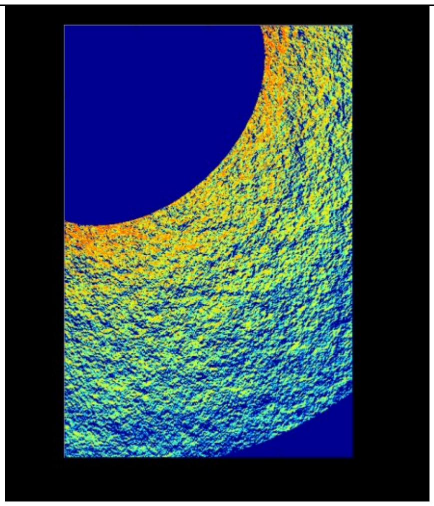

Fig. 10: Bistatic reverberation strength in dB of a fractal sandy-mud seafloor (Tx at [0m 100m], Rx at [100m 10m]).

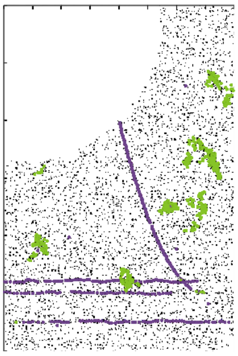

Fig. 11: Accumulated view of the HISP filter's output.

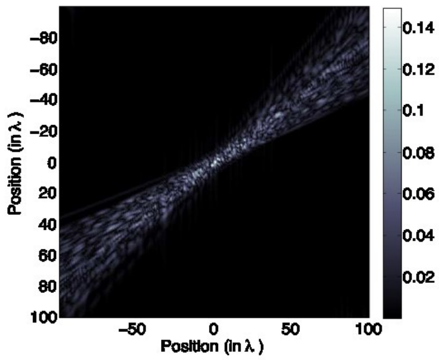

Fig. 12 MIMO image of a point scatterer target with an accuracy of the transducer locations of 5λ.

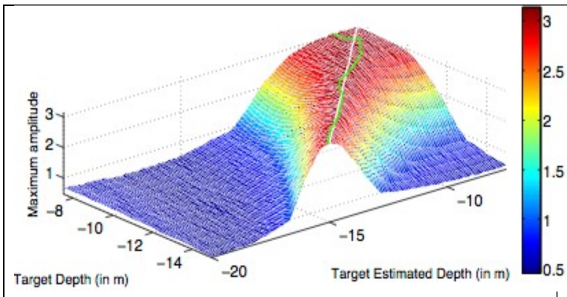

Fig. 13: Autofocus function to maximise, white curve: real target depth, green curve: estimated target depth.

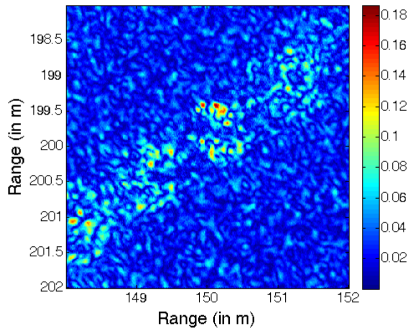

Fig. 14: MIMO image of a moving target with random transmitting delay introducing an average of 4/5 target displacement.

Subsequent to the MIMO simulator, we developed a series of algorithms and applications based on autofocus algorithms [P29]. The MIMO images are computed using the multi static back-projection algorithm, which is a variant of the bi-static back-projection algorithm developed by the SAR community. For the multi-static scenario the continuous integration along ellipses is replaced by a finite sum in which each term corresponds to one transmitter/receiver pair contribution. Using the point scatterer hypothesis, the autofocus algorithms are then based on energy maximisation.

At first, we developed a calibration algorithm based on genetic algorithm. MIMO and SAS systems share the same principles for synthetic aperture imaging and in both cases back projection algorithms can be applied to produce high-resolution images. So even for static MIMO systems the relative positions between transmitters and receivers has to be known with a sub-wavelength precision (Figure 12). Using autofocus algorithms, we demonstrated that it is possible to calibrate a MIMO sonar system within λ/10 accuracy. Such calibrated system can then be used coherently and then produce super-resolution images.

We also developed two variants of the multi-static back projection algorithm to estimate the depth of a mid-water target, its speed and orientation. To estimate the target depth we move the projection plane through the water column. A simple autofocus algorithm estimates the target depth with great accuracy (Figure 13). The main idea of speed and direction autofocus is to introduce a defocus for moving objects and then estimate the speed and direction parameter $\boldsymbol{v}$ which will re-focus the target (Figure 14).

In parallel we developed an automatic acoustical tracker by taking advantages of the multiple transmitters and receivers of MIMO systems. The automatic acoustical tracker is an extension of the time reversal mirror introduced by Prada. The main idea is to defocus the time reversed signal accordingly to the maximum speed of the target. The time reversal mirror equation then becomes:

$$\boldsymbol{E}_{2N} = \left[\prod_{2n=2}^{2N} \mathbf{G}\mathbf{K}^*_{2n-1}\mathbf{G}^*\mathbf{K}_{2n-2}\right]\mathbf{E}_0$$

where **G** is the defocus function.

Moving from harbour surveillance scenarios to coastal area survey or anti-submarine warfare (ASW) applications, we developed a new method for acoustic wave propagation in shallow water environment [P31]. For higher frequencies (above 1 kHz), the most popular method for wave propagation in shallow water is based on Ray theory and geometrical acoustics. In the ideal case scenario of a constant sound velocity profile and perfectly flat interfaces for the surface and the seafloor, an elegant solution is derived from the Mirror theorem: source images are easily geo-metrically computed by successive symmetries of the source itself (Figure 15). We proposed an extension to the Mirror theorem to take into account any interface geometry or sound velocity variation (horizontally or vertically) by solving the eikonal equation using the Fast Marching algorithm (Figure 16).

Accessing all the MIMO signals requires the orthogonality of the waveforms. As purely orthogonal waveforms do not exist, different approaches were developed to minimise the waveform cross-correlation. Such methods include TDMA (time division multiple access) where waveforms share the same frequency band, but at different times, FDMA (frequency division multiple access) where waveforms occupy different frequencies at the same time, or CDMA (code division multiple access) where waveforms share the same frequencies at the same time. Radar design and electronics impose a certain number of constraints on the waveform design. One of the most restrictive is due to the non-linear amplifiers used for such systems and imposes to the radar waveform a constant amplitude. Although the constant amplitude requirement maximises the pulse energy, it drastically reduces the degrees of freedom. The radar community are able to find efficient solutions to manipulate the signal phase including phase shift. For active sonar systems, pulse emission is the result of piezo-electric material excitation via linear amplifiers. Sonar systems are then not constrained to pulses with constant amplitude. The transducers however cannot handle drastic phase shift and phased coded waveforms may be extremely distorted through PZT transducers. We proposed a CDMA strategy which fits the requirements of wideband large MIMO sonar systems: (i) wideband width covered by every pulses, (ii) ’good’ auto- and cross-correlation functions, (iii) possibility to generate a large number of orthogonal waveforms, (iv) waveforms with smooth phase transition and (v) waveforms with relative constant amplitude. To fulfil the requirements previously stated, we propose to build the MIMO sonar waveforms using interlaced micro-chirp series (IMCS) with constant bandwidth. The waveform is the summation of two concatenations of micro-chirps series with equal duration. Each μ-chirp is chosen randomly between the sub-bands with a random up or down chirp structure. The randomised up or down structure minimised the cross- correlation as well as the sidelobes in the auto-correlation function. Figure 15 draws an example of IMCS waveform structure in the time-frequency plane. Blue and red segments represent respectively the μ-chirp structure of the first and second μ-chirp series.

The ExRaise trials took place at the acoustic calibration barge in the Bedford Bassin at Halifax, Canada between the 4th and 15th May 2015. The trials, led by DRDC Atlantic, were a collaboration between DRDC, Heriot-Watt University and Dstl. The main objective of the experiments was to investigate the acoustic response of mine-like objects (aluminium spherical shells with different contents) using a large band of frequency. The DRDC LF system (ExRaise) is using the microfine Tx and have a 10 - 50 kHz band width. The BioSonar system covers 30 - 140 kHz. Finally the Kraken SAS system was able to provide high resolution SAS images of the scene. The three sonars were mounted on a horizontal beam as drawn in figure 16. The beam attached to a pole was lowed at a depth of 18 m. The trolley was driving back and forth along a bridge with a span of around 8 m. The trolley was travelling at a speed of roughly 10 cm.s−1. One of the main interests of these trials is the plurality of waveforms used during these experiments. In particular we tested the IMCS MIMO waveforms to assess their performances in real environment. Figure 17 and 18 show preliminary results comparing traditional chirp pulses with the novel IMCS waveforms. Although chirps offer better image contrast (around 1.5dB), the IMCS approach allow multiple waveforms to be present in the environment at the same time.

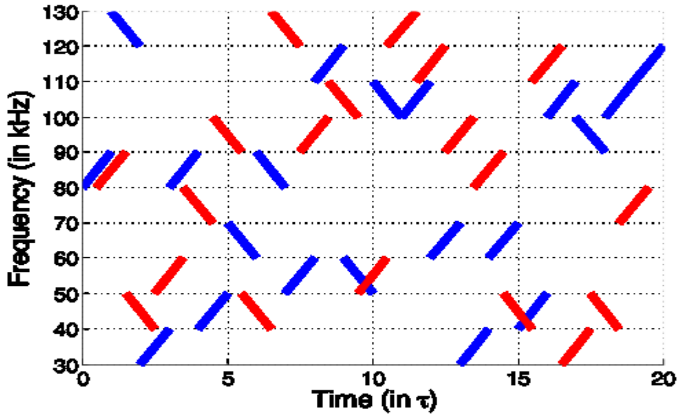

Figure. 15: Example of an IMCS waveform structure in the time-frequency domain. The blue segments represent the µ-chirps of the first series, the red ones represent the second series.

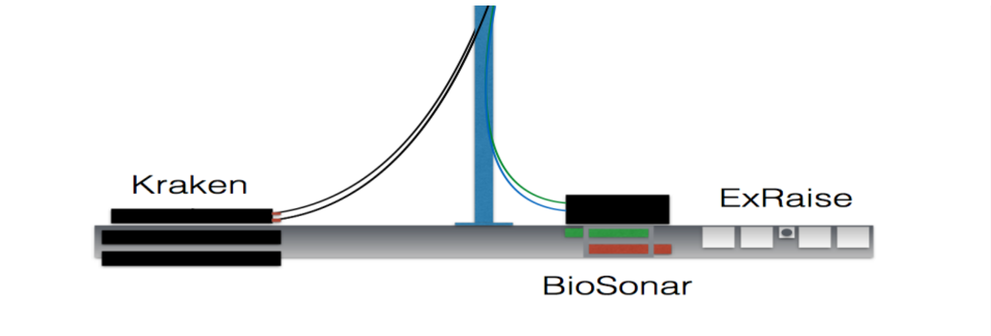

Figure 16: Sonars placement on the supporting beam.

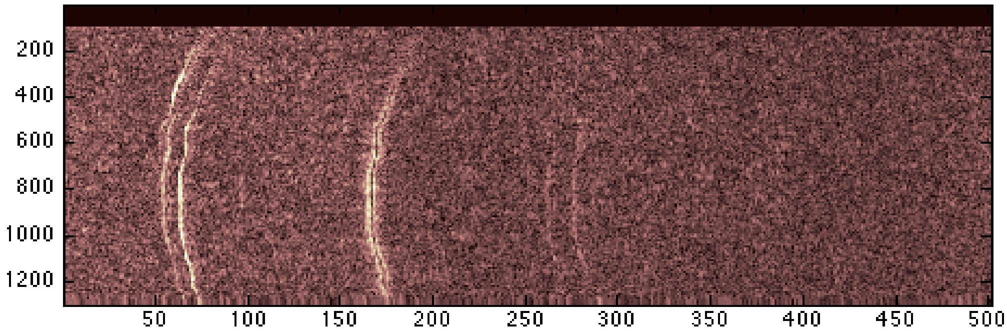

Figure 17: Target echo image using traditional up-chirp.

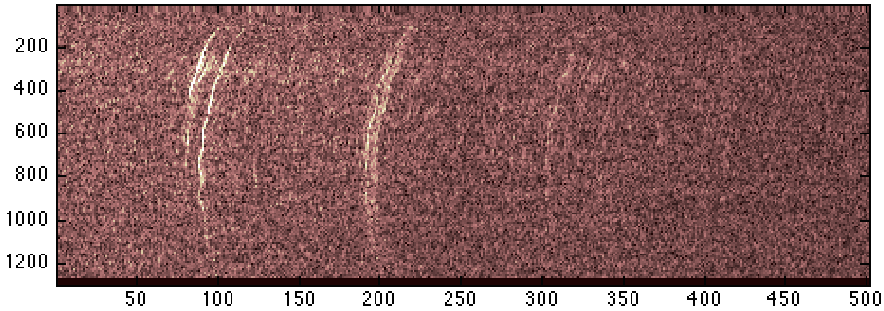

Figure 18: Target echo image using IMCS waveform

At the start of his PhD studies, Jose Franco focused on accelerating Emmanuel Delande's DISP filter by finding segments that are suitable for parallelisation and implementing them on a Graphics Processing Unit. This was motivated by the possibility of applying this filter on the maritime data that will be provided later in the year by Maritime Collaborative Enterprise (MarCE), acquired on a harbour scenario from multiple sensors mounted on a research vessel, including sonar, radar, Electronic Support Measures (ESM), GPS and Automatic Identification System (AIS) and so perform multiple-target tracking, sensor registration and calibration, and even motion pattern based target classification.

The filter was divided on a per-track part and a per-hypothesis part. The first efforts at parallelisation were directed to the per-track part, which involved using the GPU to compute Kalman Filter prediction and update for many different tracks at once, which involved implementing linear algebra routines to operate on many matrices simultaneously, in addition to coming up with efficient data structures that exploit the memory bandwith acceleration features of the graphics card. It was decided that the per-hypothesis part would be implemented on a fast serial language rather than in parallel.

Jose Franco has attended the provided courses at the supercomputing centre at The University of Edinburgh (EPCC) on GPU and parallel computing on ARCHER in order to receive training to better perform these tasks. An interesting by-product of the focus on accelerating multiple object tracking filters was a collaboration on an article related to GPU acceleration to perform calibration with the PHD filter [P35], published in ICCAIS 2014.

In the last few months, Jose dedicated time to extract meaningful measurements from the MarCE data in order to test the capabilities of the DISP filter for tracking multiple objects in a harbour scenario. The data was decoded from the binary format in which it was received. So far, basic beamforming has been applied to the (passive) sonar data, which has yielded meaningful directional information, and a basic boat detector has been prepared to work with the video data from the boat-mounted cameras. Work now remains on pre-processing the radar data, which has been decoded into a readable state, and to feed all these measurements into a filter.

The first year of Puneet Chhabra’s PhD programme was completed in October 2014, and Puneet is now proceeding to develop work on full waveform LiDAR processing. To date, he has been using simulated data from an assumed airborne platform surveying natural land terrain. While this allows progress, we do need to collect and analyse real data. A number of proposals have been made for new data collection and we also made a request in August 2014 to access Dstl data, subject to clearance and status.

For man-made object detection, DLC is critical in the underwater and aerial domains where image based techniques are widespread (Side-Scan, SAS, Light Detection and Ranging (LiDAR)). This is the focus of WP3.3, being investigated by Puneet Chhabra, a PhD student. To date, the focus has been on multi or hyper-spectral LiDAR data and on the aerial domain, e.g. the detection of vehicular targets under foliage, which requires the detection of a weak multispectral LiDAR signal in high clutter, along with the associated formation and management of the complex point clouds to give surface data that can be interpreted against target models. Of course, such techniques can be equally applied to multispectral data in the maritime domain, to detect and classify mines or other incongruous objects. As a companion study, we have also been looking at the adoption of similar techniques underwater using sidescan sonar data.

Figure 19. Illustrative output of the LiDAR processing system

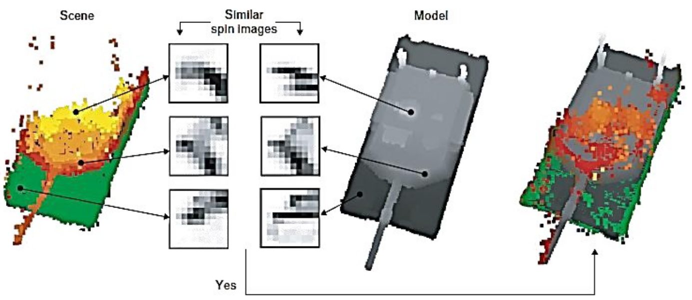

Addressing the problem of anomaly and target detection in clutter, Puneet has investigated a two-stage process, in which full waveform hyperspectral LiDAR was first screened for anomalies using wide footprint data, then attention was focussed on detailed point clouds for DLC. The developed framework detected spectral and time anomalies in the full wave multi-spectral data, using a learnt subspace (dictionary) and co-occurring local-patterns. This work is continuing to ensure a more robust and efficient algorithm to find such anomalies. We have also looked at how pattern such as spin images can be used to classify targets in monochromatic point cloud data, as shown for example in Figure 19.

Interactions

We have had regular meetings with Dstl, our two main points of contact are Gary Heald for MCM related work and Gary Davies for ASW and sonobuoys related activities. Beside the various themed meetings, we maintain regular contact with our points of contact. We have had access to valuable data and information from Dstl (MUSCLE raw SAS data, passive array data from harbour environment, technical report on ASW, sonobuoys deployment scenarios) which has helped considerably in crafting realistic scenarios for MIMO sonar systems focusing our effort in applications of defence interest.

We have spoken to Richard Green from Dstl about the availability of airborne, registered LiDAR and hyperspectral data. However, we have not been able to secure that data, so we are looking at the possibility of using data from an aerial survey conducted as part of our concurrent work on forest canopy sensing (using Riegl and Optech instruments). There also a concurrent funded project on underwater LiDAR using super-continuum sourced LiDAR for 3D imaging (Dstl contact, Richard Hollins). Using the photon counting LiDAR equipment, we plan to conduct some multi-spectral LiDAR trials in our own tank before the end of 2015.

We will work collaboratively with WP1 to explore Sparse Sensing techniques for SAS image formation and study if/how recent advances in SAR processing can benefit SAS, especially in the context of 3D SAR/SAS where we have real data available to perform experiments.

A new line of investigation in Wideband Sonar (WBS) for target detection in clutter will be opened. WBS enables the tracker to acquire the response of targets and clutter across a wide range of frequencies. Spectral analysis can be used to discriminate targets of similar shape based on their acoustic response, especially using low frequencies where the internal resonances of targets differ depending on material and structure. This analysis has strong relationships with multi-spectral LiDAR and we expect joint work in this area. We already have a PhD currently working on these aspects (Dstl, but not UDRC funded) and we will collaborate with her on these aspects.

WP4 Context Driven Behaviour Monitoring and Anomaly Detection

Research Leader: Neil Robertson

Academics: James Hopgood

Research Associate: Rolf Baxter

PhD Student: Alessandro Borgia

This work package investigates the identification and classification of behaviours as normal or abnormal, safe or threatening, from an irregular and often heterogeneous sensor network. Although some general principles may be applied, there is unlikely to be a unified approach that can be applied across different sensor domains; for example, tracking moving 'blobs' in radar or sonar data is quite distinct in dimension, mathematical and applied treatment from tracking human subjects in CCTV image data. In this work package we focus specifically on the problems of using electro-optic (video, IR) and audible data to monitor behaviour and detect anomalies, therefore addressing a number of specified areas of interest and challenges.

Our approach to anomaly detection is based on the premise that better models of normality are required for more complex anomalies to be detected. As such this work package aims to identify techniques for improving behaviour models by using spatio-temporal context, and will develop novel techniques for better predicting future target intent (both short-term and long-term). Prediction errors will indicate that the learnt behaviour models do not match observed behaviour, and are indicative of a potential anomaly.

Outcomes

Our work on intentional priors has been well received and has been published in IEEE Signal Processing Letters [P33], in addition to 2 conferences (P35, P36). This work proposed the use of intentional priors for providing better predictions of target motion, and showed that person head-pose is well correlated with target motion. By developing a novel (Deep Belief Network) head-pose classifier we showed that important contextual information can be extracted and used to improve predictions about future target behaviour (e.g. during occlusion).

To extend this work further we have been developing a novel clustering algorithm for learning target pattern-of-life. This algorithm will form the basis of a second intentional prior and will form a component of our unified contextual tracking and anomaly detection algorithm. An outcome of the work so far has been a (winning) submission to the Dstl Wide Area Motion Imagery (WAMI) anomaly detection challenge. Based on our pattern-of-life models we were able to identify numerous anomalies within the Wright Patterson Air Force Base dataset.

We have also had success in the Dstl temporal anomaly detection challenge in which our algorithm was able to identify several challenging anomalies not revealed by other entries. Follow-on work has since been discussed with Dstl, with a call for tenders pending.

Additionally, our Pattern Recognition journal article has now been accepted for publication following a minor revision [P32]. This work addressed behaviour recognition and anomaly detection within a video surveillance application, and was a continuation of work initially sponsored by MoD under the competition of Ideas initiative.

Progress

WP4.1 initially focused on joint audio-video tracking in complex occluded environments. By integrating and fusing audio and video signals at different levels of abstractions we showed that the multimodal tracker is more reliable than single modality systems, tolerating large occlusions and cross-talk within enclosed environments [P34, P41, P43].

Extending this work we then proposed a novel method for detecting and localising the dominant speaker in an enclosed scenario [P37, P40]. Audio source position estimates are computed by a novel stochastic region contraction (SRC) audio search algorithm for accurate speaker localisation. This audio search algorithm is aided by available video information (stochastic region contraction with height estimation (SRC-HE)) which estimates head heights over the whole scene and gives a speed improvement of 56% over SRC. In contrast to the current state-of-the-art, we enlarged and extended the field of view and considered cluttered scenarios (e.g. multiple non-stationary moving speakers). In this work, we learn Mel-frequency cepstral coefficients (MFCC) coefficients and correlate them to the optical flow. By exploiting the audio and video signals to estimate the position of the actual speaker, the visual search-space is narrowed and reduces the probability of incurring wrong voice-to-pixel region associations. Comparing our work with the state-of-the-art on real datasets showed a 36% improvement in localisation precision.

More recently, we have predominantly focused on improving target behaviour models by using spatial context within the video domain. Specifically, we developed a framework for detecting short-term behavioural priors (termed intentional priors), and integrated these with target motion models to improve target tracking under the premise that better models of expected behaviour are required for anomaly detection. We demonstrated the approach by developing novel head-pose based priors for short-range video surveillance (e.g. concourse protection). Based on the statistical observations that people tend to look in the intended direction of travel, we integrated head-pose information into a person tracker and showed that target tracking performance could be improved via a head-pose prior. After reporting early results in Sensor Signal Processing for Defence [P35], we extended this work by developing a novel deep belief network head-pose classifier and demonstrated improved pedestrian target tracking on two public video datasets. This work has now been published in IEEE Signal Processing Letters [P33].

We have also demonstrated the utility of contextual information in crowded surveillance scenes using a novel unsupervised context-aware process. We found that in a crowded scene the application of Mutual Information based social context permitted the ability to prevent self-justifying groups, and to propagate anomalies through a social network. The strength of our contextual features is demonstrated by the detection of subtly abnormal behaviours, which otherwise remain indistinguishable from normal behaviour [P42].

Our most recent work has seen a transition from short-term behavioural priors (e.g. head-pose) to longer-term priors (e.g. pattern-of-life). To model spatial context we have been developing a clustering algorithm inspired by the Piciarelli model for clustering target trajectories. This has been evaluated against the Dstl provided Wright Patterson Air-force Base (WPAFB) dataset, and additionally to our own dataset of target tracks gathered using GPS loggers. By integrating spatial context into the model learning process, models of normal target behaviour can be identified over large geographical regions (e.g. wide area motion imagery). Not only does this allow target behaviour to be anticipated using pattern-of-life model estimates, but anomalous behaviour can also be identified. Anomaly detection based on our spatial context model has achieved a false positive rate (FPR) of 0.0076 and a true positive rate (TPR) of 1 on the WPAFB dataset, and a FPR of 0.21 and TPR of 0.66 on our GPS track dataset. Furthermore, our algorithm maintains on-line real-time modelling and anomaly detection over large areas, attributed not exhibited by previous techniques.

We have also been using kernel density estimation techniques to model temporal context by learning the activity distributions associated with different spatial locations. The efficient representation of these distributions is a critical step for long duration missions and one aspect of our work has focused on techniques that remain efficient as the number of observations increases.

Most recently, we have been developing an automated approach for configuring the anomaly detection threshold based on the target detector false-detection rate. Initial results on simulated data have been positive and - pending full integration into our temporal model - should provide a robust and parameter light solution for spatio-temporal anomaly detection in wide areas.

Future Direction

Through discussions with Dstl it has been agreed that pattern-of-life (POL) learning will remain the focus in the coming year while audio-video fusion will remain on-hold. We will continue to develop our WAMI POL algorithms and begin work to integrate learnt POL as an alternative ‘intentional prior’ within a unified target tracking framework. As with our head-pose ‘prior’, determining the contextual cues for invoking/switching between priors will be essential, and we will be developing a generic framework for identifying and switching between different priors. As our models of ‘normal’ behaviour are improved, we will return to the anomaly detection aspect by extending initial work based on monitoring tracker innovation.

WP5 Network Enabled Sensor Management

Research Leader: Daniel Clark

Academics: Yvan Petillot, Mike Davies

Research Associates: Emmanuel Delande, Jeremie Housinneau

PhD student: Jose Franco, Alexey Narykov

Defence applications now routinely deploy a network of sensors (mobile or fixed) with heterogeneous and complementary sensing capabilities to achieve a particular set of goals. Advances in sensor technologies have led to the emergence of a large number of degrees of freedom in sensing devices (waveform selection and adaptation for radar, sampling schemes in compressed sensing, etc.). More importantly, these degrees of freedom can be controlled by software, providing agility in their real time operations

There is an existing growing body of work in this area, mostly in the US with applications such as waveform agile radar. Sensor management is currently seen as a decision process and is based on Markov Decision Processes (MDP) and their partially observable version (POMDPs). The output of the process is a policy, which determines the optimal sensors allocation and configurations at each time step, based on previous measurements and expected reward. The problem can, in theory, be solved within a Bayesian filtering paradigm that models the system as a whole to satisfy operational constraints (time, exclusion zones, and energy) and achieve operational objectives (target detection, behaviour analysis, situation awareness). While the definition of the reward is obviously goal-specific, generic reward (entropy minimization for instance) has proven effective in a large category of problems. There are, however, key obstacles to the real deployment of these algorithms. First, the computational complexity is currently prohibitive for any multi-stage planning (more than one step ahead), second, the models assume perfect communications between sensing entities, a centralized coordinator for the network and unlimited bandwidth. In this work package, we propose to study the sensor management problem in the context of a networked set of mobile assets and sensors.

In WP5.5, hierarchical sensor management for target tracking, the determination of higher-level operational decisions involves a hierarchy of problems. At the sensor level, this involves the development of observation models for detecting, characterizing and classifying targets from a variety of different sensors as well as adapting signalling to targets and local clutter (WP3.3). At the object/environment level, there is a need to be able to distinguish high-priority targets from background environment, whilst building an increased awareness of the environment. At the decision-level, methods are required to control and synchronise the sensors to maintain focus on targets of interest. Integrating this information in a unified manner will enable the information at the operator level to be reduced to avoid an overload of information. This work package will focus on high-level statistical and information-theoretic multi-sensor fusion and sensor management. The underlying objective will be to develop a coherent and unified framework for integrating compatible Bayesian estimation, control, and sensing techniques.

For WP5.2, computationally tractable solutions, these will be developed using recent ideas in compressed sensing, finite set statistics (see WP3.1) and convex optimization. We will investigate suboptimal solutions where performance can be predicted and controlled. We will study how the problem space can be reduced by working at a higher level of abstraction in the DPDC loop (Direct-

Capture-Process-Disseminate) whilst delegating the low-level sensor management to individual asset or groups of assets. Hardware implementation of solutions and links with computation resource management will be tackled into WP6.3.

WP 5.3 multi-objective sensor management will focus on developing robust methods to network and assets failure and taking into account multiple goals when the number of degrees of freedom in the system can provide redundancy that can be exploited to perform multiple objectives simultaneously. Recent developments in control theory, where objectives can be prioritized and extra degree of freedoms used to achieve secondary objectives could be usefully exploited.

Outcomes

The key outcome is an underpinning mathematical framework for sensor management that enables the detection, estimation, localisation and tracking of objects in cluttered environments ranging from heterogeneous sensors [P52, P56], decision-based sensor control based on population activity [P53, P20, P54, P55], calibration of different sensors onto a common reference frame [R17, P58], and estimation of moving objects from a moving sensor platform [R18, P59, P35]. In addition is the ability to classify different types of objects based on dynamics or observation characteristics [P27].

A novel multi-target detection/tracking algorithm, the DISP filter [P52], was designed and implemented in order to fuse the information of heterogeneous sensors in a challenging environment (clutter, missed-detection, unknown number of targets, etc.) into a unified and coherent framework. Fully probabilistic in nature, it is compatible with both sensor management approaches explored so far (see next paragraph). The DISP filter has been exploited on a single-sensor space surveillance scenario for the monitoring of space debris [P60] in collaboration with Dstl's Space and Strategic Systems group, Purdue University, and the US Air Force.

Two tools of very different nature have been developed for the design of sensor management policies oriented toward the detection and tracking of multiple targets:

a) The second-order regional statistics for population-based filters [P23] provide the estimated number of targets, with given uncertainty as maintained by the tracking filter, in any region of the surveillance space selected by the operator. It is an inexpensive statistical tool that provides a broad picture of the surveillance scene in order to identify the regions where uncertainties about local target activity are the highest. Available for any multi-object filter that can be derived from the Finite Set Statistics framework [R14], we have explored their application for several pre-existing filters [P23, P54, P55] and illustrated their exploitation for the PHD filter on a sensor management scenario [P62].

b) An information-theoretic approach [P53] has been developed in the context of closed-loop sensor management for which the reward function can emphasizes specific targets and/or specific regions of the surveillance space. More complex in nature and more computationally challenging, it provides ground for a more customized sensor policy according to the operator's needs (some targets may not be worth focusing on, some regions should be covered in priority by available sensing resources, etc.). It was designed for any multi-object filter than can be derived from the novel estimation framework for stochastic populations [R15], including the DISP filter [P52], but its principle could be extended to the traditional track-based approaches used in industry, most notably the MHT filter [R16].

Progress

The general motivation behind the work produced in WP5 is the development of a statistical estimation framework for target detection/tracking and sensor management that is sufficiently general to encompass the problems of detecting, localising and tracking of multiple objects in cluttered environments from different types of sensors, and the principled design of sensor management policies from a probabilistic and unified description of all the system uncertainties (multi-target configuration, target behaviour, sensor observation process, multi-sensor calibration, etc.).

The design and implementation of the DISP filter [P52] from the recent estimation framework for stochastic populations [R15] forms the backbone of the developments in WP5 towards the construction and exploitation of a unified multi-target/multi-sensor estimation framework, while the HISP filter [R15] has also been exploited to permit classification of different types of objects based on dynamics or observation characteristics in a maritime surveillance scenario [P27].

In the context of space situational awareness, for which Dstl's Space and Strategic Systems group aims at unifying all the UK ground-based assets focussed on space surveillance (radars, lidars, cameras, telescopes, etc.) into a coherent framework enabling these heterogeneous observations to a single multi-object estimation algorithm. This has proven to be a challenging and extremely motivating research context from which an ongoing fruitful collaboration with Dstl, Purdue University and the US Air Force has formed. As a first step, the DISP filter was illustrated on a single-sensor space surveillance scenario involving a Doppler radar in a conference paper [P60], which will be followed by a journal submission to the AIAA Journal of Guidance, Control, and Dynamics before the end of the year.

Ongoing developments within D. Clark's research group aim at incorporating a wide range of sensors to the DISP framework, covering the variety of the UK ground-based assets. The modelling of telescopes and their incorporation in the unified framework has recently been completed by J. Franco, and should be the topic of conference paper to the 2016 IEEE Aerospace Conference (Big Sky, U.S.A.). In parallel the processing of real observation data from two UK facilities, a radar from the STFC Chilbolton facility (UK) and a range of cameras from the Dstl facilities in New Zealand, are ongoing. Finally, inspired from a novel method of jointly calibrating different sensors onto a common reference frame developed in the group [R17], a procedure has been developed to jointly estimate satellite trajectories and the jittering/drift of a ground-based telescope from the sequence of frames produced by the sensor. It should be the topic of another conference paper to the 2016 IEEE Aerospace Conference.

Following D. Clark's invitation as a visiting professor to the department of Astrodynamics of the University of Colorado Boulder for a month-long summer course on multi-target tracking, we have initiated a collaboration focussed on the design of a novel CPHD filter augmented with target spawning, in order to cope with specific space scenarios where targets (space debris or man-made objects) disintegrate into smaller units. The main results are secured on ArXiV [P63], and it will be the topic of journal submission before the end of the year.

The development of second-order regional statistics for population-based filters [P23, P56] has been followed by their implementation for a various number of existing filters [P23, P28, P54, P55] and illustrated on a sensor management scenario for the PHD filter in [P62]. This concept was originally developed for any multi-object filter derived from the FiSST framework [R14], and it is also directly applicable for any filtering solution derived from the more recent estimation framework for stochastic population [R15], such as the DISP filter [P52]. An inexpensive statistical tool describing the estimation of the level of target activity, with given uncertainty, in any region of the surveillance region, it can be deployed in multi-sensor networks to ensure that optimal decisions can be taken to aid global situational awareness in the surveillance region.

The development of an information-theoretic approach for closed-loop sensor management [P53] was motivated by the availability of specific information on identified targets (i.e. "tracks") in both traditional track-based solutions such as the MHT filter [R16] and recent estimators for stochastic populations such as the DISP filter [P52], that per construction cannot be exploited by population-based sensor management approaches such as the regional statistics (see paragraph above). The solution in [P53] proposes a sensor policy that can focus on acquiring information on arbitrary targets and/or arbitrary regions of the surveillance space as defined by the operator, and it could be augmented to incorporate sensor-related costs associated to a specific decision of the sensor management (energy consumption, sensor inavailability while in transition to a new sensing profile, etc.). It is, however, much more computationally challenging than the second-order regional statistics; the ongoing GPU-based parallel implementation of the DISP filter by J. Franco will be exploited to address this challenge.

A novel method of jointly calibrating different sensors onto a common reference frame has been developed [R17]. The method was then applied in camera networks to be the first method of calibrating cameras from multiple moving objects [P58]. These methods were also developed for jointly tracking multiple targets and stationary landmarks from a moving platform [R18] that has been used in an underwater vehicle application [P61], and in super-resolution microscopy [P59]. The work underpins and collaborates with the work in WP2 on distributed localisation and tracking. A computationally tractable implementation has been developed and demonstrated with a GPU [P61]. These methods will be developed further for submarine command and control under a new project with Dstl that started in January 2015.

Future Direction

Alongside the development of the new estimation framework and the design of the DISP filter, we have initiated a fruitful collaboration between academia (Heriot-Watt University, Purdue University), the branch of Dstl engaged in space surveillance activities, and UK ground-based facilities for space surveillance (STFC Chilbolton Observatory, Dstl facilities in New Zealand). A key objective relevant to UK space and security interests is to fuse the information collected from the network of UK ground-based surveillance assets spread all over the Earth into a coherent picture of the objects (satellites, man-made or natural debris) orbiting around Earth, coordinate the sensing assets and cue the sensors for specific surveillance missions in order to maintain an up-to-date catalogue of orbiting objects.

Significant progress has been made on the development of the multi-target detection/tracking algorithm for simulated orbital targets through the DISP filter, and on the exploitation on real data collected from radar and telescope facilities. In light of this collaboration that has attracted significant interest from both academia and industry, and provides a practical challenge driving the development of novel solutions within the thematic challenges of WP5, we propose to set the coordination of the network of UK surveillance assets to produce a unified estimation framework for the space situational awareness problem as the main challenge that will motivate the future developments of WP5.

An investigation on the use of second-order population statistics in different regions for sensor management is currently being conducted. There are two strands to this research:

a) Information-theoretic sensor management. Reference [P53] developed the underpinning mathematics for making decisions on sensor control based on populations of tracks presented at SSPD 2014. Though the methodology is a development of the advances in the theoretical framework in multi-object estimation made within the group, the approach can be used within existing track-based infrastructure and is thus relevant to current architectures. Current work is on the development of algorithmic techniques based on Gaussian mixture approaches and investigation of realistic cost functions for sensor control. Specific scenarios will be identified with discussions with Dstl and tested in realistic signal processing environments.

b) Regional variance based sensor control. References [P23, P28, P54, P55, P62] develop a new statistic that has previously been unavailable to operators - the variance in the number of targets in different regions of the state space. This was originally developed in the context of PHD/CPHD filters, though [P23] uses a model for track-based infrastructure that is widely deployed and hence the approach can be used within existing track-based multi-target trackers and is thus relevant to current architectures. With the general formulation, however, such architectures can be extended to include additional surveillance monitoring in regions where the knowledge of activity is much lower. Current work is on the investigation of population models that permit high variance in the number of objects, so that we can deal with parts of the surveillance region where little is known in a principled and holistic way. This will inform the sensor management strategies developed.

WP6 Efficient Computation of Complex Signal Processing Algorithms

Research Leader: John Thompson

Academics: Andrew Wallace, Neil Robertson

Research Associates: Calum Blair

PhD Student: Saurav Sthapit

The overall aim of this part of the programme is to allow the deployment of complex signal processing algorithms which are relevant to the networked battlespace concept in a wider variety of devices and environments. A key part of this involves understanding the relationship between the algorithms we wish to run, and the constraints imposed by the processing and communications hardware on which they will be implemented. This work package directly addresses Dstl technical challenge 29 (Reducing Size, Weight and Power Requirements Through Efficient Processing).

We look at this relationship in several modalities: Synthetic Aperture Radar, Synthetic Aperture Sonar, visible-light imagery and ambient Radio Frequency (RF) signals. As part of WP6.1 (Efficient Parallelization of Sensing Processing) we cover three Dstl technical challenges: c5 (3D SAR Processing), c27 (Accreditable machine learning or data-driven techniques) and c29, while WP6.2 (Implementation of Distributed Signal Processing Algorithms) concentrates on c29.

Our SAR work focuses on the hardware acceleration of state-of-the-art algorithms generated in WP1.2, in particular the acceleration of compressive sensing (CS)-based synthetic aperture radar image formation algorithms. Many of the state-of-the-art SAR image formation algorithms for systems with sensor constraints are based on techniques from the CS literature. Such systems include three-dimensional (challenge 5), ground moving target indicator (GMTI) and low-frequency (LF) SAR. For CS based image formation algorithms to be computationally feasible, fast forward and backward algorithms for the SAR observation model must be available. Computationally efficient algorithms for these forward and backward operations (the fast re-projection and back-projection algorithms) were developed in UDRC phase 1. One of the goals of WP6.1 was to investigate the use of Graphics Processing Units (GPUs) to exploit parallelism in these algorithms to further reduce the image formation times of CS based algorithms.

WP6.1 Efficient Parallelization of Sensing Processing also aims to improve accreditable machine learning (c27). For any sensing modality, our aim is to generate fast, low-power detections which are accurate and have a reliable level of confidence associated with each. If an unusual data sample is presented to the detector which is considerably different from the training data, this should be reflected in a low confidence in the classified output. Current state of the art detection algorithms are optimised to produce accurate detections at the expense of reliable confidence scores, which often results in classifiers which are massively overconfident in their predictions. WP6.1 investigates techniques to improve this score generation while preserving the accurate performance of current state of the art classifiers.

Our first application is pedestrian detection in images and videos. This is an active area of vision research and classifiers developed for this problem perform well when retrained on other vision tasks (object detection and surveillance). As a demonstration that these techniques generalise well, we apply them to SAS imagery used for detection of mine-like objects. As the dataset contains multiple shape classes, we can train classifiers to separate several classes from the background and each other. At test time we also introduce previously unseen classes, and aim to detect these and flag them as needing further review using our uncertainty detector.

WP6.2 Implementation of Distributed Signal Processing Algorithms improves the performance of a distributed sensing network. We are using the example of localization of one or more hostile RF emitters, using a sensor network where communication between nodes is heavily constrained. While the ability to accurately localize an emitter is governed by physical limitations such as geometric dilution of precision (GDOP), our aim is to achieve accuracy as close as possible to the ideal (described by the CRLB or Cramer-Rao Lower Bound). We are simulating this problem in a MATLAB environment, although we will consider the possibility of a real-world localization experiment as our algorithms mature.

These overall goals of the work package described above also mesh well with Saurav Sthapit's PhD research. The broad aims for his PhD are to model the power consumption of different components of mobile phones and networks of phones. This will help in studying the feasibility of running Computer Vision, Person re-identification and Anomaly Detection algorithms on several mobile phones as opposed to running them on a single device. The trade-off between on-board computation and distributed computing will be studied to determine the best possible use cases to design and run the algorithms.

Outcomes

A conference paper [P64] showed various approaches for improving the reliability of state-of-the-art classification algorithms; Figure 20 shows the output from a state-of-the-art classification algorithm with extremely confident false positives. (b)- (d): Postprocessing algorithms reduce confidence of false positive detections. This led to further work with Seebyte as part of a 2014 MarCE task, where the reliability and uncertainty techniques we used were applied to SAS imagery. This allowed us to classify existing objects and also to identify the presence of objects which we had not trained the detector to look for [P66]; see Figure 21.

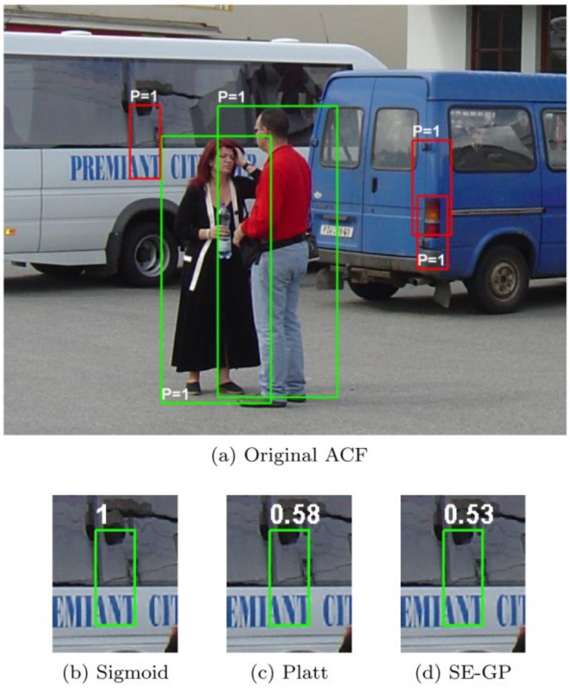

Figure 20: improving the reliability of detected windows in a pedestrian detection algorithm. The original algorithm is overconfident in its detections.

This problem of "classification with confidence" allows human operators or autonomous systems to be able to rely more on the certainty of the detections that algorithms generate. If a detection is certain, it may need no further investigation, while if an initial detection is flagged as 'uncertain', then another sensor pass or further processing of that region may be required.

A follow-up paper was also presented [P68]. This focused more closely on the computational aspects of introspective classification, and produced an accelerated Graphical Processor Unit (GPU) implementation of Gaussian Process Classifiers, demonstrating a 3.7x speedup for a pedestrian detection application.

[P65, P67] were based on a power consumption comparison of Field Programmable Gate Arrays (FPGA) and GPU. In a system with SWaP constraints, power consumption, detection speed and algorithm accuracy are all considered important. This research investigates on-the-fly selection of processor accuracy to improve speed when anomalous behaviour is spotted, while saving power in routine scenarios and show the reductions in power consumption available if these techniques are applied.

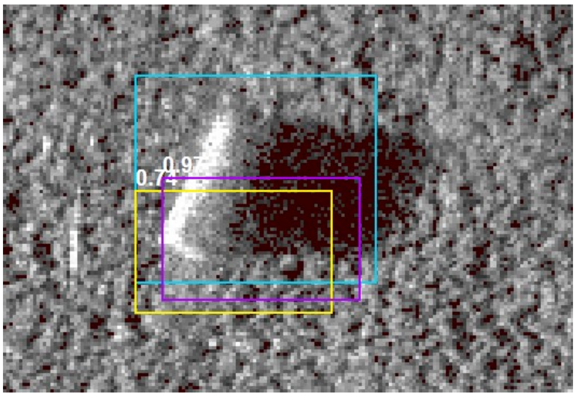

Figure 21: Detecting an unseen and untrained object (the cylinder with blue box) using uncertainty information from trained cone and wedge shapes (purple box). A false positive from the wedge (yellow) is also shown.

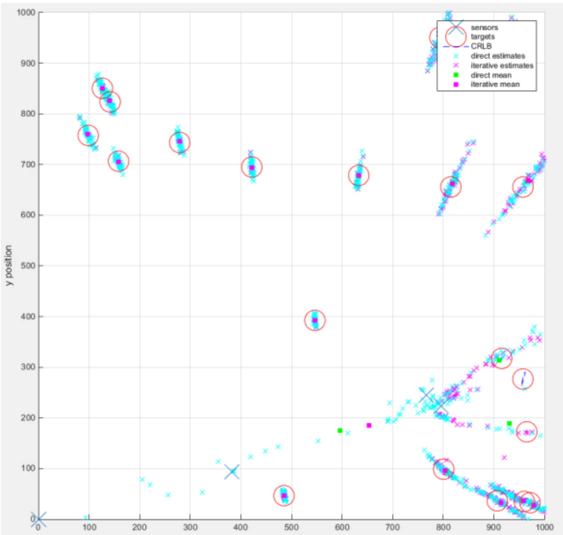

Figure 22: Simulated localisation of emitters (pink circles) from a network of receivers (blue crosses)

A simulator for evaluating distributed localization algorithms has been developed. This allows us to test methods for detecting and localizing RF emitters in scenarios where the processing power and communication bandwidth available to the detection network is constrained. We are using this to improve the accuracy of localization algorithms used on devices running in a mobile ad-hoc network (MANET). We have concentrated on algorithms based on Time Difference of Arrival (TDOA) and in each case we are comparing these to the theoretical localisation accuracy achievable as measured by the Cramer-Rao Lower Bound (CRLB). See Figure 22.

The parallelisation techniques used to speed up SAR backprojection were written up in [P8].

A paper comparing algorithms for pedestrian re-identification was published [P69]. This compared three algorithms for feature extraction and comparison, evaluating them on the basis of re-identification accuracy, volume of data to be transferred between nodes, and computational complexity.

Progress

Considering SAR first [P8, P9], the progress has involved modifying the fast algorithms so they are suitable for GPUs. Fast (re/back)-projection algorithms have two main components, recursive decimation/upsampling and standard (re/back)-projection. These two components cannot be done in parallel. However, within these components, parallelism may be exploited. The following is a list of modifications which have been made to the algorithms to allow them to exploit a parallel processing platform.

- The standard (re/back)-projection algorithms were modified to make use of recent advances in non-uniform FFTs (NFFT) to reduce the memory usage. Memory usage is highly constrained on GPUs.
- The loop ordering of the computational elements in the standard (re/back)-projection algorithms were modified so they are now independent. The independent elements can be used as parallel threads on a GPU.
- The recursive decimation/upsampling structure in the algorithms has been replaced by a loop structure. Loop structures are more suitable for parallel processing than recursive structures.
- Investigation into the choice of decimation/upsampling kernels based on speed and accuracy.

Work on 3D SAR is now being taken forward as part of WP1.

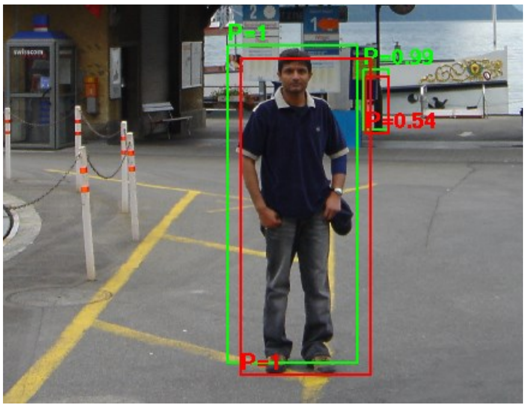

Figure 23: improving the reliability of detected windows in a pedestrian detection algorithm. The original algorithm (shown in green) is always overconfident in its detections, but an improved algorithm (red) allows better identification of uncertain detections.

Other work in WP6.1 has dealt with the 'classification with confidence' problem. We have completed and published a study into the ability of classifiers to provide confidence measures which are more reliable and hence make machine learning algorithms more accreditable. This will act towards improving operator trust in complex tasks involving detection and classification algorithms. It may also improve the performance of tracking, anomaly detection and other higher-level inference algorithms which rely on the confidence an object detector exhibits when stating that a particular object is of a specified class.

We considered the problem of pedestrian detection, a well-studied task in computer vision. We were able to show improvements in reliability in existing state-of-the-art detectors at the cost of a limited increase in computational time [P64]. Existing algorithms tend to be overconfident, producing predictions which are associated with a confidence score of 100% even for marginal samples. Data-driven techniques for remapping these to a more realistic confidence score were applied. Figure 23 shows that a high confidence in true positive object detections can be maintained while reducing the confidence measures assigned to false detections. In addition, algorithms such as Gaussian Process classifiers (GPCs) were trained, as these had been shown previously to produce more realistic confidence measures. However, the computational complexity of GPCs is orders of magnitude higher than that of Adaboost and can take minutes per frame when classifying pedestrians.

A paper describing a parallelised implementation of this on GPU was presented [P68] and the code has been made available for download. This shows a 3.7x speed up over existing optimised versions.

These techniques for improving reliability are not specific to pedestrian detection and can be applied to other modalities. We have also looked at MSTAR, a dataset containing SAR imagery of multiple vehicle types. We could classify vehicles into types with increased reliability, and also assign higher uncertainty to vehicles which did not belong to one of the training classes.

However, a more detailed investigation was carried out using similar techniques applied to SAS datasets (Colossus and Catharsis MUSCLE-SAS data). These contain multiple object types and various levels of background clutter. Again, we were able to improve the reliability of existing detectors and show that samples with the most uncertainty correspond to object types which the classifier has not been trained upon. We focused on a comparison of support-vector machines (SVMs) and GPCs, showing that SVMs demonstrated increased discriminative ability when detecting objects from known classes. They also proved able to indicate the potential presence of a previously unseen class. An example detection of a cone-shaped mine-like object using a SVM classifier is shown in Figure 24. A conference paper describing this work has been presented [P66] and a journal paper is in preparation.

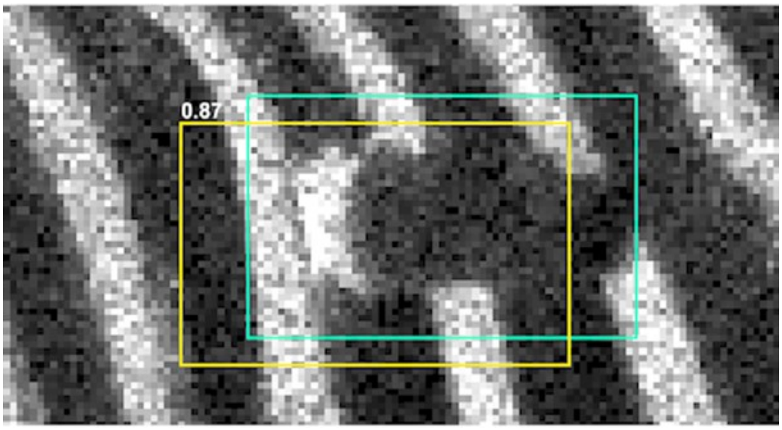

Figure 24 Detection of a mine-like object in SAS imagery using a SVM classifier

In other work relating to improving SWaP, a paper detailing the results of a power consumption study between FPGA and GPU for object detection in video was presented at a computer vision applications conference [P65], and a longer journal version has been accepted for publication [P67]. This quantifies the improvements available in power consumption of a video surveillance system when using a measure of anomalous behaviour to dynamically select which architecture to process imagery on. We switch to faster GPU processing when unusual behaviour is registered and return to lower-power FPGA processing when the observed behaviour is normal. We plan to make the source code used for this experiment available for download before publication.

We have also completed some work on WP6.2; this has involved discussions with Ben Gear at Dstl. Our distributed signal processing work has focused on the task of accurately identifying the position of one or more hostile RF emitters, using a distributed network of receivers. We are constrained both by the processing power of the receivers and the low bandwidth available for communications between them (in practice, the bandwidth is liable to dominate). Time difference of arrival-based (TDOA) algorithms appear to be the most effective method so far, but we are also using the received signal strength as an initial step to allow spectrum characterisation and node selection before performing final localisation. Physical constraints also affect the accuracy achievable, with the Cramer-Rao Bound (CRB) giving the minimal achievable error in any noisy system; if this is too high to be effective we must also look at moving the position of some of the sensors or using additional sensors to improve precision. Figure 25 shows our simulator in action; rather than sharing all TDOA signals received by our sensor network, we select the nodes which give the lowest CRB based on initial received signal strength information, and transmit TDOA information from them to a fusion node. This trades off localisation accuracy against available inter-node bandwidth and power consumption.

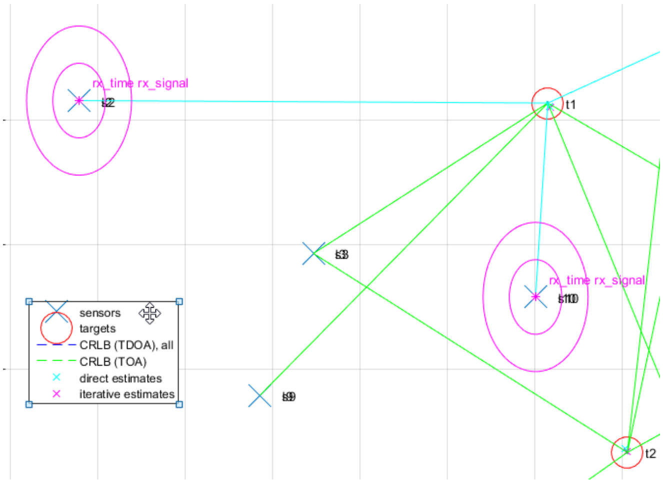

Figure 25. Distributed localisation simulator showing selected nodes transmitting received signal data to a fusion node.

Saurav Sthapit's work on smartphone platforms has concentrated on identifying suitable object detection, recognition and anomaly detection algorithms and measuring their power consumption. For this purpose, Android applications were developed and experiments were carried out to compare the energy and time required to run the algorithm locally on a smartphone vs. the energy and time required to send the raw data to the server for processing, where it was assumed that the server has no power or computing limitations. Initial results are shown below and there are no unexpected surprises in the result. The preliminary results show that simple algorithms for face detection can be run in real time on current smartphones. It consumes less than 1 Joule of energy and the calculation time is in the millisecond range. Offloading it to be processed in the cloud would not only incur higher energy usage but also suffer from high end to end delay. For more complex algorithms, such as HOG (Histogram of Oriented Gradients – a commonly used pedestrian detection algorithm) which takes tens of seconds to process, it may be wiser to offload processing to the cloud or to other peer devices. It could result in better response time and energy savings as well.

Suitable smartphone-capable anomaly detection algorithms have also been investigated. We have concentrated on face recognition algorithms ranging from Principal Component Analysis (PCA) to DeepFace, the current state of the art. We built a prototype based on PCA on MATLAB emulating server and client. The client calculates the feature vectors and probes the server for identification. When contacted, the server searches the database for any match and replies back with the answer. For the face recognition training and testing purpose, we have collected about 100,000 images. This work is still at an early stage but ongoing.

We have recently focussed on person re-identification in multiple cameras based on the appearance of the person as a whole. We have identified key algorithms and descriptors used in the field and performed some experiments to determine their suitability to run on distributed systems. Figure 27 below shows the performance of various methods for person re-identification, namely Symmetry Driven Accumulation of Local Features (SDALF), Keep It Simple and Straightforward Metric (KISSME) and Saliency matching. Use of facial features would be optional depending on the image resolution. This work was recently published in [P69].

Current work involves simulating computer vision algorithms running on a network of sensors. The main simulation objective is to find effective ways of offloading data and co-ordinating between the sensors. By simulating the time and energy consumption during computation and communication, we can apply different strategies to achieve our goals such as saving energy, running huge algorithms, maximising network lifetime etc. Figure 28 below shows a simulation of the distributed sensor network, and Figure 29 shows the maximum number of targets Google Nexus N1 can process based on an algorithm with 10 million operations per target if the smartphone is not running anything. Knowledge of these characteristics allows us to understand how processing performance can be traded off against node longevity when the number of targets increases. Finally, Figure 30 shows an output of our power consumption vs CPU load simulation.

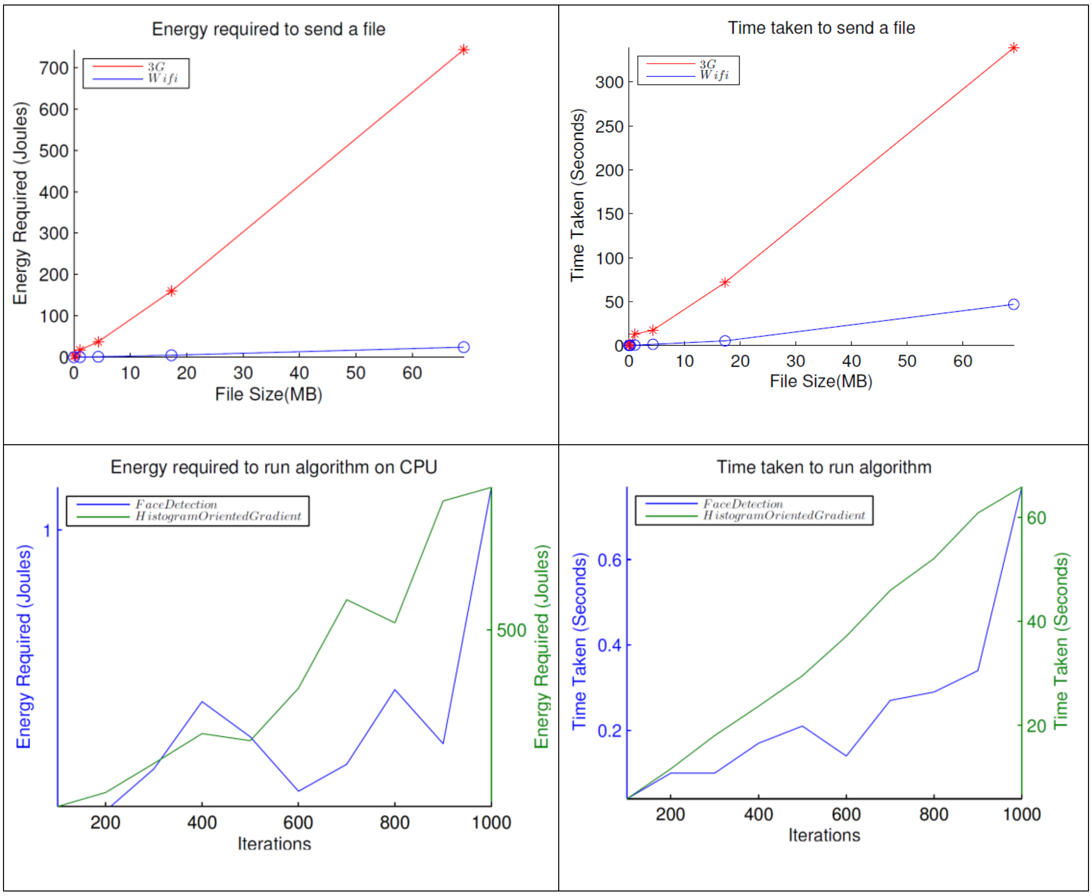

Figure 26: Runtime and energy consumption for algorithms run on mobile phone

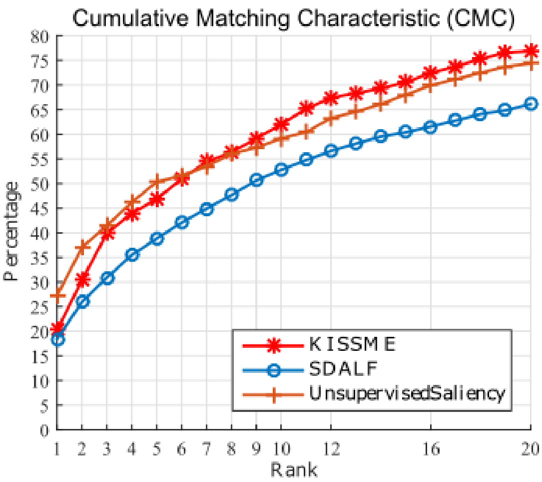

Figure 27: Performance of person re-identification algorithms as rank vs. percentage of samples correctly identified. The rank of a test sample refers to its position in the ordered list of possible matches in the dataset.

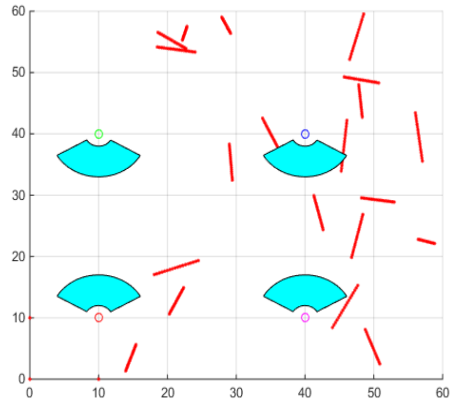

Figure 28. Simulation of a distributed network of non-overlapping cameras (blue) and randomised pedestrian tracks (red)

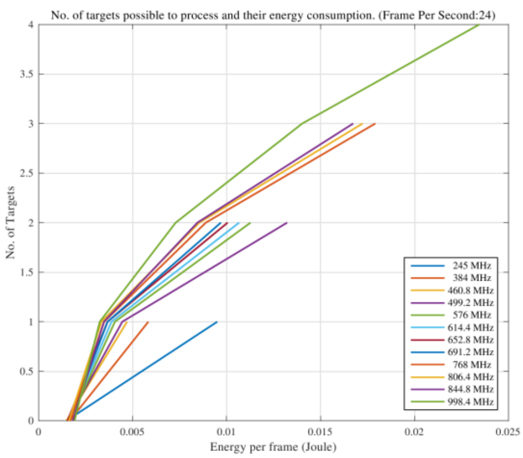

Figure 29: A power vs. performance graph of distributed target tracking. This shows the maximum number of targets a Google Nexus N1 can process based on an algorithm with 10 million operations per target if the smartphone is not running anything.

Figure 30 Sample simulator result.

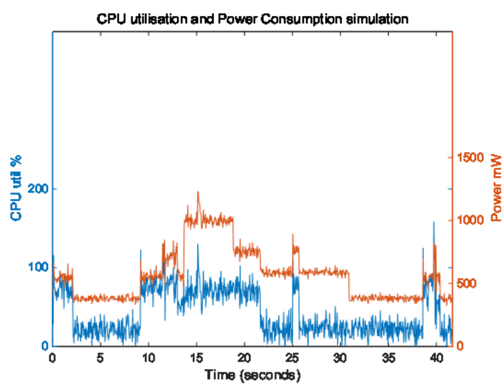

Future Direction

Most of our effort in the next year will concentrate on the implementation of distributed signal processing algorithms and we will extend our existing simulator to focus on the problem of communication constraints and remove some abstractions relating to the way signals are emitted. The software could then be published. Another consideration will be the running of a real-world localization trial to demonstrate algorithm performance.

A longer journal article will describe all our classification work from WP6.1 to date, focusing on computational efficiency and the ranking of algorithms used in terms of accreditability or reliability, leading to a comparison of current deep learning algorithms for applications including in the SAR, SAS and imaging domains.

In the last quarter of this year we will commence work on WP6.3: Algorithm & computation resource management. This has been extended from the initial proposal to run for the remaining two years. This will concentrate on scenarios where it may be infeasible to fully process all available incoming data on our available compute resources, both for computational (time) reasons, and limitations in Size, Weight and Power (SWaP). We will investigate novel approaches for attending to individual sensors or areas of sensors, or re-tasking sensors based on the results of gathered data. This will also exploit methods developed in WP6.2 work for distributed processing and resource management.

Engagement

Engagement, communication and dissemination are key to the success of the UDRC. The Edinburgh Consortium have set up a detailed strategy which has created a two way communication channel between UDRC (both Edinburgh and LSSC Consortiums) and interested stakeholders.

As part of this strategy, Edinburgh has developed and manages 2 websites and runs a series of annual events involving the signal processing community. The rise in awareness of the research carried out by the UDRC has been supported through strong marketing and has created strong connections with this community around the world. Following on from the success of the UDRC email newsletter campaign, which saw an above average open and click percentage creating a high volume of traffic to the UDRC website, we have also seen an increase in requests to join the UDRC LinkedIn group. Each of these digital applications has contributed to the successful marketing campaign of the UDRC.

The marketing campaign included generating further interest in the UDRC annual events, namely the Summer School and the SSPD Conference. The UDRC has opened a line of communication with other leading UK Universities in signal processing including Cambridge, Bristol, Manchester and Newcastle, who have agreed to promote the UDRC events and consequently are now connections for future events. We are now starting to replicate these relationships with international universities.

The URDC has also featured in articles written for the MOD Defence Contracts Bulletin magazine, Forbes magazine, The Herald and Financial Times. We are also working closely with the Knowledge Transfer Network for Sensors and Photonics and the Innovation Centre for Sensor and Imaging Systems (CENSIS) who are able to market our events and research across their networks.

The UDRC website has proved an extensive resource for informing the members of the UDRC and the wider signal processing community about the research and researchers involved in the project. The resources webpage is an excellent archive which shares relevant material of past events which has been invaluable for sharing ideas within the project and informing the new members of the UDRC community. The events section of the website has not only been useful for marketing the UDRC and SSPD events it has also been used to promote UDRC related events which has enabled connections with other projects and signal processing groups. As part of the marketing and engagement element of the UDRC, connections have been forged between other UK universities which specialise in signal processing. The website has been useful in creating these relationships by advertising other projects events and in turn the UDRC events have been promoted on their external websites. In order to rate the effectiveness of the website, it is monitored through Google Analytics which provide useful statistics giving an insight into the website users.

The Audience Analysis feature enhances our understanding of what the users expect from the site and through the recording of the key words and the sites which referred the user to the UDRC website, it is possible to understand what the target audience is using the site for, therefore enabling the Team to determine a trajectory for the online presence of the project. The results from Google Analytics can be viewed on the next page.

Google Analytics Results from April 2013 to September 2015

How many people visited the website?

- Since the launch of the new UDRC and SSPD website over 12,100 people have visited the sites.
- The websites have increased the amount of users by 150.03% from year 1 (April 13 - March 14) to year 2 April 14 - March 15).
- The page views of both websites have stayed consistent with a slight rise from 47,252 in year 1 to 47563 in year 2.

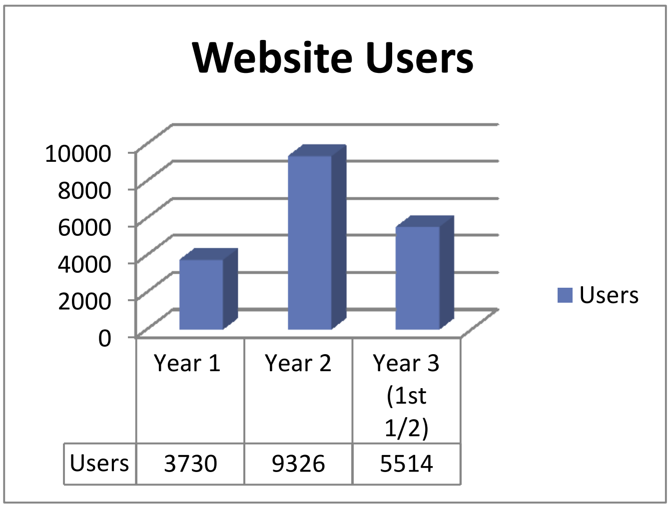

We can see that there was a significant rise in the number of users of the UDRC website, with the year 3 projection set to have an increase in website users than previous years. The google analytics also depicts the type of user actively accessing the website, such as a new or returning viewer. This statistic has fluctuated from year to year with each year of the project reporting over half of the visitors had not visited before, which could represent the response from the UDRC events and marketing strategy.

How engaged were these visitors within the website?

- Over the course of the project users spent on average 4 minutes viewing the UDRC webpages and each viewer on average viewed 6 web pages in total.
- The type of visitor to the site has also changed with the percentage of returning visitor increasing from 34.6% (Oct 13) to 57.8% (March 15).

The UDRC phase 2 has sent out three successful newsletters to all those affiliated with the UDRC and the 843 people who have subscribed to UDRC updates, informing the recipients about useful dates and links to events. The first was emailed December 2014, this winter edition of the newsletter was successfully delivered to 753 recipients and was opened 586 times. The open rate (33%) of this email newsletter was double the industry (education) standard according to Mailchimp, the application used to disseminate the newsletter. This was improved slightly in the most recent newsletter to a 34% open rate.

The top clicked links in the three newsletters were:

1. http://www.mod-udrc.org/news
2. www.sspdconference.org
3. www.mod-udrc.org/events/2015-summer-school

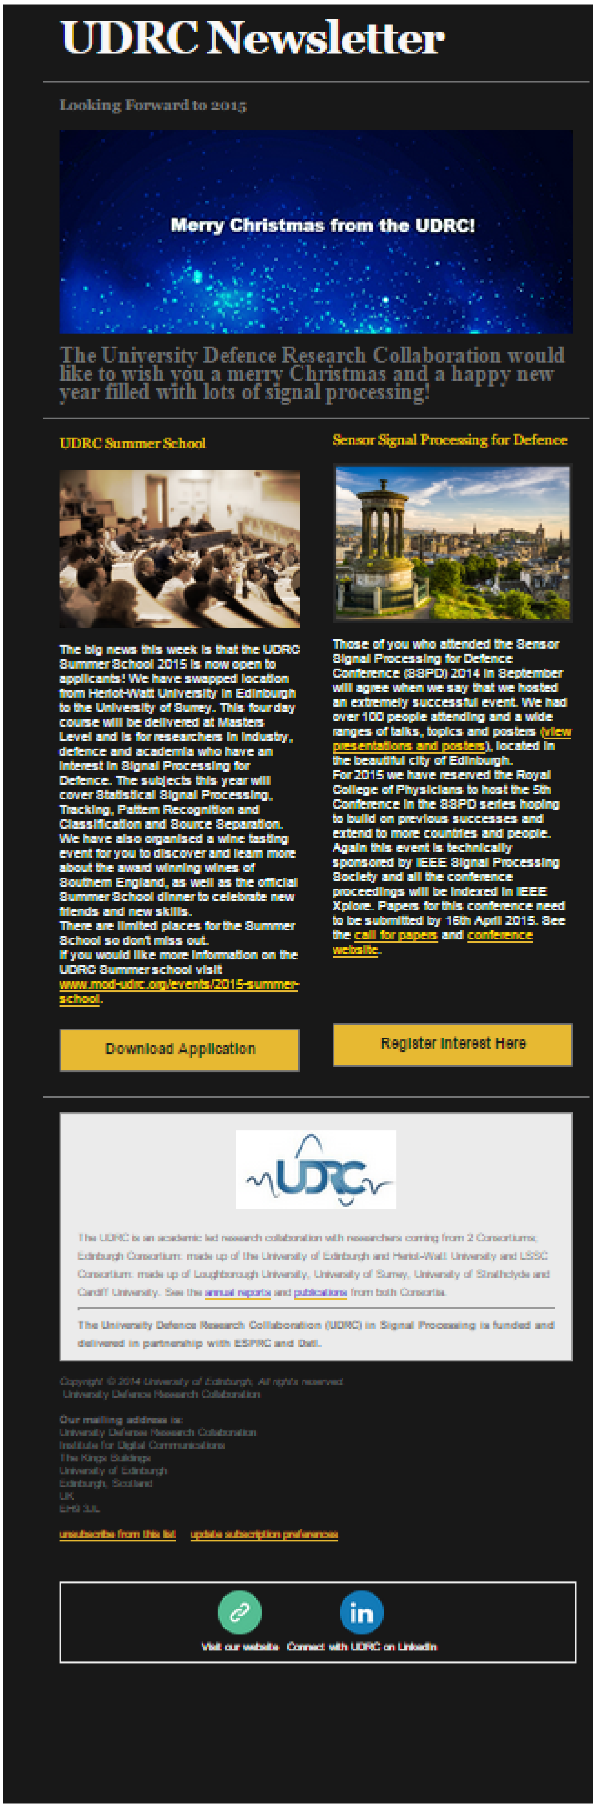
UDRC Newsletter

Looking Forward to 2015

Merry Christmas from the UDRC!

The University Defence Research Collaboration would like to wish you a merry Christmas and a happy new year filled with lots of signal processing!

UDRC Summer School

The big news this week is that the UDRC Summer School 2015 is now open to applicants! We have swapped location from Heriot-Watt University in Edinburgh to the University of Surrey. This four day course will be delivered at Masters Level and is for researchers in industry, defence and academia who have an interest in Signal Processing for Defence. The subjects this year will cover Statistical Signal Processing, Tracking, Pattern Recognition and Classification and Source Separation.
We have also organised a wine tasting event for you to discover and learn more about the award winning wines of Southern England, as well as the official Summer School dinner to celebrate new friends and new skills.
There are limited places for the Summer School so don't miss out.
If you would like more information on the UDRC Summer school visit www.mod-udrc.org/events/2015-summer-school.

Download Application

Sensor Signal Processing for Defence

Those of you who attended the Sensor Signal Processing for Defence Conference (SSPD) 2014 in September will agree when we say that we hosted an extremely successful event. We had over 100 people attending and a wide ranges of talks, topics and posters (view presentations and posters), located in the beautiful city of Edinburgh.
For 2015 we have reserved the Royal College of Physicians to host the 5th Conference in the SSPD series hoping to build on previous successes and extend to more countries and people.
Again this event is technically sponsored by IEEE Signal Processing Society and all the conference proceedings will be indexed in IEEE Xplore. Papers for this conference need to be submitted by 16th April 2015. See the call for papers and conference website.

Register Interest Here

The UDRC is an academic led research collaboration with researchers coming from 2 Consortiums; Edinburgh Consortium: made up of the University of Edinburgh and Heriot-Watt University and LSSC Consortium: made up of Loughborough University, University of Surrey, University of Strathclyde and Cardiff University. See the annual reports and publications from both Consortia.

The University Defence Research Collaboration (UDRC) in Signal Processing is funded and delivered in partnership with ESPRC and Dstl.

Copyright © 2014 University of Edinburgh, All rights reserved.
University Defence Research Collaboration

Our mailing address is:
University Defence Research Collaboration
Institute for Digital Communications
The Kings Buildings
University of Edinburgh
Edinburgh, Scotland
UK
EH9 3JL

unsubscribe from this list update subscription preferences

Visit our website Connect with UDRC on LinkedIn

The implementation of the UDRC seasonal newsletter helps to gain an insight into the international awareness of the UDRC. The top locations the newsletters were opened:

UK - 60%

USA - 22.2%

Canada - 6%

Germany - 1.9%

China - 1.1%%

<u>LinkedIN</u>

The social media platform LinkedIn has proved very successful in not only marketing UDRC and SSPD events but allowing UDRC contacts to ask questions, share information and build relationships with the signal processing community. The UDRC group has made connections with other signal processing LinkedIn groups, with comments and advice being exchanged in this social network. To date the UDRC group has 84 members and has been very useful in reaching out to an international audience, resulting in delegates travelling from across Europe and Asia to attend UDRC events.

Events

Events are a key part in the success of the research project and as part of the coordination process a number of meetings, workshops and events have been organised and managed by Edinburgh Consortium in partnership with LSSC Consortium and Dstl.

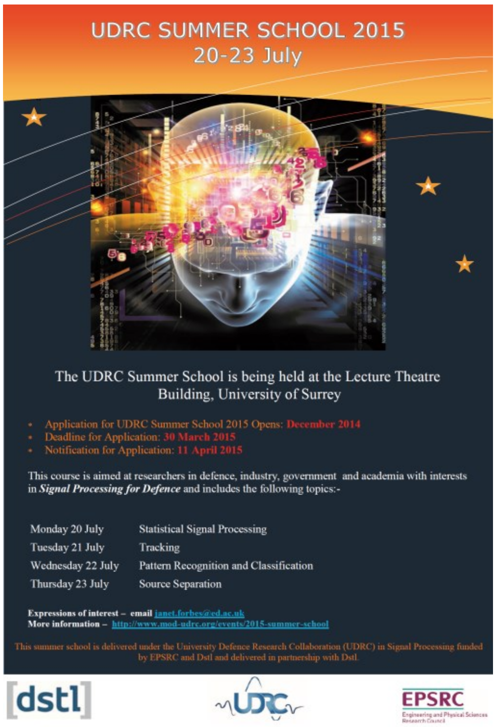

Summer School

Since 2013 the annual UDRC Summer School has been organised by the Edinburgh Consortium and over a 4 day period covers a range of topics on signal processing for defence. It is an informative week with researchers from both consortiums involved in teaching and fostering research in the field of signal processing. Participants have to apply before they are allowed to take part and from over 80 applicants ranging from Europe, Asia, Africa and America, only 60 are invited to attend the classes. Furthermore, a social program including a Summer School dinner and whisky and/or wine tasting is also offered.

We have successfully organised four Summer Schools; two in 2013 on finite set statistics and one in 2014 on target detection and tracking, compressed sensing, anomaly detection and source separation and one in 2015. The latest summer school took place at the University of Surrey and covered the subjects of statistical signal processing, tracking, pattern recognition and classification and source separation. This year's attendees came from South Korea, India, UK, Czech Republic, Germany, Italy, Pakistan, Canada, China, Iran, Russia, Iraq, Nigeria, Turkey and Bulgaria.

Industrial Day

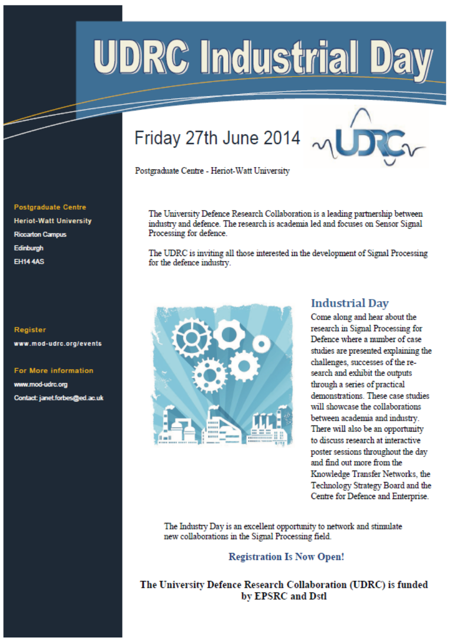

Industrial Days are all about showcasing the UDRC research and knowledge transfer between industry and academia. We held an industrial day in 2014 which was a great success and had just under 80 participants attending with speakers and exhibits from industry, academia, the Centre for Defence and Enterprise, the Technology Strategy Board and the Knowledge Transfer Network. Our next industrial day is planned for 2016 where we aim to attract a multi sector audience in signal processing with a view to discussing and developing ideas in differing applications such as medical or scientific monitoring.

Themed Meetings

Themed meetings are meetings which enable the experts from both Consortiums, Dstl and industrial partners to get together and discuss topics in more detail. As well as discussing ideas and potential collaborations there is always an afternoon session which presents two Dstl challenges in the area of discussion. This is discussed in a small workshop and the attendees then have a chance to submit their solutions to Dstl. Winners of these challenges are awarded a prize at subsequent meetings.

Since 2013 we have organised five successful themed meetings on the following topics; source separation, NIS ITA joint meeting on communication and networks and distributed/multi-sensor/source processing, anomaly detection, autonomous systems and signal processing and MIMO and radar signal processing.

These meetings, as well as developing research ideas and collaborations, have been a huge success in terms of developing potential research areas which have materialised out of the Dstl challenges. One challenge from Dstl on spectral deconvolution has initiated contract work on the design of a Raman spectral deconvolution model using results from a Dstl Raman data set. Another challenge on Ground Penetrating Radar led to Mehrdad Yaghoobi receiving a Dstl prize for the best entry. In one of the latest challenges on temporal anomaly detection, Dstl has asked for a longer contract of work to be developed on transient device detection from noisy time-series data.

Knowledge Transfer Meetings

Knowledge Transfer Meetings are an annual initiative to promote and raise the profile of knowledge transfer and innovation within Dstl and the academic researchers. The purpose of these knowledge transfer meetings is to catalyse and facilitate communication and innovation. The previous two meetings involved the Consortiums presenting their work with ensuing discussion on Dstl research needs and ways to work together to achieve them. A third meeting that took place in April 2015 had a slightly different format. The layout consisted of focused application based workshops led by Dstl on various applications such as advanced uses of phased array sonar and space, the final frontier. These stimulated discussion on how the current research can be developed and restructured to address these applications. Academics were also able to clarify how specific MOD challenges can be addressed for example in the workshop on advanced uses of SAR, specific problems were discussed and a solution was proposed by the academics for their challenges.

Events from April 2013 to December 201

Event	Date	Description
UDRC Summer School	Jun-13	2 x Finite Set Statistics summer schools (USA/UK)
Source Separation and Sparsity	Oct-13	Themed meeting with academics and industry presenters and Dstl challenges
UDRC Launch	Dec-13	Research presented with Dstl and industry keynotes
NISITA Joint meeting	Jan-14	Themed Meeting on how ITA and UDRC can work together
KTM	Mar-14	Knowledge Transfer Event at Dstl with academia and defence
Uncertainty and Anomaly Detection	May-14	Themed meeting with academics and industry presenters and Dstl challenges
UDRC Summer School	Jun-14	Successful Summer school held in Edinburgh
UDRC Industrial Day	Jun-14	Case studies of joint work with industry and academia
SSPD 2014	Sep-14	Well attended event with industry, military and academic perspectives
Open Source Tracking meeting	Sep-14	Meeting to take forward open source tracking framework
Open Source Tracking meeting	Oct-14	Follow-up meeting
Autonomous Systems and Signal Processing	Nov-14	Themed meeting with academics and industry presenters and Dstl challenges
Sonar meeting	Nov-14	Meeting with Nick Goddard and academics
Dstl/Academic workshop	Dec-14	Dstl workshop looking at research needs within Defence
KTM	Apr-15	Knowledge Transfer Event at Dstl with academia and defence
MIMO and radar signal processing	May-15	Themed meeting with academics and industry presenters and Dstl challenges
UDRC Summer School	Jul-15	70 people signed up for University of Surrey Summer School
SSPD 2015	Sep-15	100 people attended
Hardware and implementation themed meeting	Nov-15	Themed meeting with academics and industry presenters and Dstl challenges

Signal Processing for Defence (SSPD) Annual Conference

The Edinburgh Consortium has organised a series of conferences beginning with the UDRC Launch in December 2013 leading on to the SSPD conference in 2014. The Launch was a fantastic opportunity to trigger the assemblage of the signal processing community with 124 delegates attending. Each year the annual conference goes from strength to strength becoming an established event in the signal processing calendar, attracting invitations from other conferences to collaborate. The SSPD conference in 2014 attracted 100 delegates and 39 papers were submitted. A technical programme committee was drawn up from technical experts from across the globe. IEEE Signal Processing Society co-sponsored the conference and the publications were indexed in IEEE Xplore. We also attracted an internationally recognised keynote speaker, Randy Moses to talk on his radar research. The video of this keynote can be found on the conference website.

We asked for SSPD conference feedback and the response was very positive. Selected comments, and suggestions for the future are below.

- Keynote speaker is terrific
- The industry session was very good
- Military session also good standard
- Would like more time for questions with the military advisors- this was the most unique opportunity of the conference.
- Have less but longer presentations allowing more time for questions
- Perhaps brainstorming breakouts or workshop session

For the SSPD 2015 conference, we attracted 100 delegates and the submission of 48 papers of which 32 were accepted. We extended the invitation to the Technical Programme Committee to include more members from the international arena. As well as technical co-sponsorship with the IEEE Signal Processing Society, we have also secured additional support from the IEEE Aerospace and Electronic Systems Society (AESS) who circulated our call for papers to all their members. The conference has also been marketed through our LinkedIn campaigns, UDRC newsletter, CENSIS, KTN, ICASSP and the KTN Intelligent Imaging Programme event where we promoted the UDRC, SSPD and Summer School at our exhibition stand.

In planning for SSPD 2016, we agreed with the International Society for Optics and Photonics (SPIE) to hold our event and their SPIE Security and Defence conference close to each other in date and location. The SSPD 2016 conference will therefore take place in Edinburgh on the 22nd and 23rd September 2016, a day before the start of the SPIE conference. This partnering should allow for an increase in attendees to SSPD2016 as people may like to attend both conferences.

In investigating potential partners for SSPD, we have also held discussions with DASP to explore the possibility of alignment. We arranged a series of telephone meetings and the outcome of these was that DASP would establish if they were able to travel to the UK as travel permission is currently difficult to obtain for Australian and American defence organisations. The organisers of the DASP conference were also invited to sit on the Technical Committee for SSPD 2015 which they have accepted.

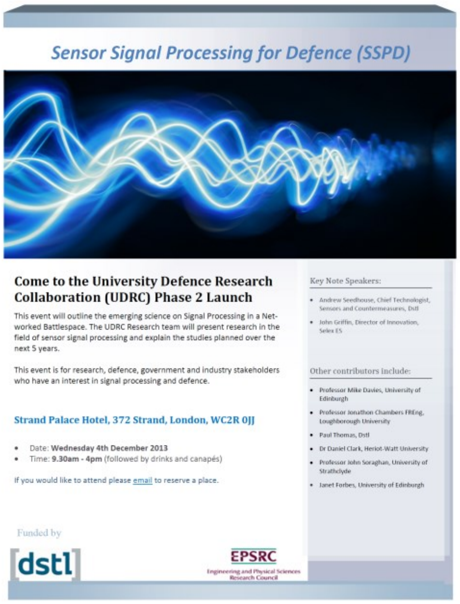
Sensor Signal Processing for Defence (SSPD)
Come to the University Defence Research Collaboration (UDRC) Phase 2 Launch
This event will outline the emerging science on Signal Processing in a Net-worked Battlespace. The UDRC Research team will present research in the field of sensor signal processing and explain the studies planned over the next 5 years.
This event is for research, defence, government and industry stakeholders who have an interest in signal processing and defence.
Strand Palace Hotel, 372 Strand, London, WC2R 0JJ
Date: Wednesday 4th December 2013
Time: 9.30am - 4pm (followed by drinks and canapés)
If you would like to attend please email to reserve a place.
Key Note Speakers:
Other contributors include:
Funded by
dstl
EPSRC


Sensor Signal Processing for Defence Conference (SSPD 2015)
Royal College of Physicians
9 Queen Street
Edinburgh
EH2 1JQ
September 9 - 10, 2015
Call for papers: www.sspdconference.org
Important Dates:
Submission of Papers:
16th April 2015
Paper Acceptance:
11th June 2015
Final version of Paper Due:
8th July 2015
Paper Submission:
Technical sponsorship is provided by the IEEE Signal Processing Society and the IEEE Aerospace and Electronic Systems Society and proceedings will be submitted to the Xplore Digital Library.
- Array Signal Processing
- Image Processing
- Radar, Sonar and Acoustic
- Multimodal Signal Processing
- Multi-Target Tracking
- Signal Acquisition and Sensor Management
- Multiple-input and multiple-output (MIMO)
- Data Fusion
- Source Separation
- Anomaly Detection
- Distributed Signal Processing
- Low Size Weight & Power Solutions
- Target Detection and Identification
- Electro-Optic Sensing
dstl
AESS
EPSRC


Sensor Signal Processing for Defence
Conference 2014
Important Dates:
Submission of Paper Deadline:
4th April 2014
Notification of Paper Acceptance
4th June 2014
Final version of Paper Due:
8th July 2014
Conference Dates:
8th and 9th September 2014
Technical sponsorship is provided by the IEEE Signal Processing Society and the proceedings will be indexed in IEEE XPLORE.
John McIntyre Conference Centre, Pollock Halls, University of Edinburgh
Topics include, but are not limited to, the following:
dstl
EPSRC


Sensor Signal Processing for Defence Conference
SSPD 2016
22 - 23 September
Royal College of Surgeons
Edinburgh
Important Dates:
Submission of Papers: 22nd April 2016
Notification of Paper Acceptance: 24th June 2016
Final version of Paper Due: 22nd July 2016
The Sensor Signal Processing for Defence Conference is organised by the University Defence Research Collaboration (UDRC) in Signal Processing. SSPD 2016 aims to bring together researchers from academia, industry and government organisations interested in Signal Processing for Defence.
Call for papers: www.sspdconference.org
Papers are solicited from the following areas:-
- Array Signal Processing
- Image Processing
- Radar, Sonar and Acoustic
- Multimodal Signal Processing
- Multi-Target Tracking
- Signal Acquisition and Sensor Management
- Multiple-input and multiple-output (MIMO)
- Data Fusion
- Source Separation
- Anomaly Detection
- Distributed Signal Processing
- Low Size Weight & Power Solutions
- Target Detection and Identification
- Electro-Optic Sensing
All submitted papers will be peer reviewed. Accepted papers will be published in conference proceedings which will be archived in IEEE Xplore. Sponsored by Dstl and EPSRC.
Expressions of interest – email janet.forbes@ed.ac.uk
dstl
EPSRC
AESS

Publications

[P1] A method of analyzing radio-frequency signals using sub-Nyquist sampling. UK priority filing patent application: 1309783.7, priority filing date: 31/05/2013.

This patent sets out our proposed framework for low SWAP sub-Nyquist sampling. A simplified acquisition, detection and reconstruction/analysis system for the detection and analysis of multiple radar signals within a wideband spectral sensing scenario. The solution uses a smart compression sensing algorithm, with a multi-coset sampling strategy in combination with the computation of some efficient Time Frequency transforms.

[P2] A Low-complexity Sub-Nyquist Sampling System for Wideband Radar ESM Receivers, M Yaghoobi, M Lexa, F Millioz and M Davies, ICASSP, Florence, Italy, May 2014.

This paper introduces a new framework for multichannel sampling of ESM signals, each with a rate much lower than Nyquist. The main contribution here is on introducing a low-complexity reconstruction algorithm for the radar pulse reconstruction.

Abstract: The problem of efficient sampling of wideband Radar signals for Electronic Support Measures (ESM) is investigated in this paper. Wideband radio frequency sampling generally needs a sampling rate at least twice the maximum frequency of the signal, i.e. Nyquist rate, which is generally very high. However, when the signal is highly structured, like wideband Radar signals, we can use the fact that signals do not occupy the whole spectrum and instead, there exists a parsimonious structure in the time-frequency domain. Here, we use this fact and introduce a novel low complexity sampling system, which has a recovery guarantee, assuming that received RF signals follow a particular structure. The proposed technique is inspired by the compressive sampling of sparse signals and it uses a multi-coset sampling setting, however it does not involve a computationally expensive reconstruction step. We call this here Low-Complexity Multi-Coset (LoCoMC) sampling technique. Simulation results, show that the proposed sub-Nyquist sampling technique works well in simulated ES scenarios.

[P3] A Computationally Efficient Multi-coset Wideband Radar ESM Receiver, M Yaghoobi, M Davies, NATO Specialist Meeting on Compressed Sensing for RADAR/SAR, Tallinn, Estonia, May 2014.

This paper investigates a comprehensive comparison between proposed method with the state of art industrial approach to Radar ESM. The proposed framework has a continuous monitoring of the spectrum, in the contrast with the RFW which observe each band at a time. The processing gain of the RFW would then be limited for the lack of continuous monitoring of the spectrum.

Abstract: The problem of efficient sampling of wideband Radar signals for Electronic Support Measures (ESM) using a parallel sampling structure is investigated in this paper. Wideband radio frequency sampling, which is a necessary component of modern Radar surveillance systems, generally needs a sampling rate at least twice the maximum frequency of the signal, i.e. Nyquist rate, which is generally very high. However, when the signal is highly structured, like Radar signals, we can use the fact that signals do not occupy the whole spectrum and instead, there exists a parsimonious structure

in the time-frequency domain. Here, we characterise a novel low-complexity sampling system with a recovery guarantee, assuming that received RF signals have a particular structure. The proposed technique is inspired by the compressive sampling (CS) of sparse signals and it uses a multi-coset sampling setting, while it does not involve a computationally expensive reconstruction step. The new framework is comparable with current rapid frequency sweeping technique, while it continuously monitors the spectrum. It is thus affected from short pulse misidentification. On the other hand, like other CS based sub-Nyquist sampling techniques, it suffers from noise folding effect.

[P4] An Efficient Implementation of the Low-Complexity Multi-Coset Sub-Nyquist Wideband Radar Electronic Surveillance, M. Yaghoobi, B. Mulgrew and M. E. Davies, SSPD 2014, September 2014.

A parallel sampling structure for the efficient sampling of wideband radar signals is investigated in this paper. A new implementation technique of [P2] is here introduced to further reduce the computational cost.

Abstract: The problem of efficient sampling of wideband radar signals for Electronic Surveillance (ES) using a parallel sampling structure will be investigated in this paper. Wideband radio frequency sampling, which is a necessary component of the modern digital radar surveillance systems, needs a sampling rate at least twice the maximum frequency of signals, i.e. Nyquist rate, which is generally very high. Designing an analog to digital converter which works with such a high sampling rate is difficult and expensive. The standard wideband ES receivers use the rapidly swept superheterodyne technique, which selects a subband of the spectrum at a time, while iterating through the whole spectrum sequentially. Such a technique does not explore the underlying structure of input RF signals. When the signal is sparsely structured, we can use the fact that signals do not occupy the whole spectrum. There indeed exists a parsimonious structure in the time-frequency domain in radar ES signals. We here use a recently introduced low-complexity sampling system, called LoCoMC [R1], which is inspired by the compressive sampling (CS) of sparse signals and it uses the multi-coset sampling structure, while it does not involve a computationally expensive reconstruction step. A new implementation technique is here introduced, which further reduces the computational cost of the reconstruction algorithm by combining two filters, while improving the accuracy by implicitly implementing an infinite length filter.

We also describe the rapidly swept superheterodyne receiver and compare it with the LoCoMC algorithm. In a contrast to the former technique, LoCoMC continuously monitors the spectrum, which makes it much more robust in the short pulse detection.

[P5] Sparsity-Based Autofocus Techniques for Under-sampled Synthetic Aperture Radar, S. Kelly, M.Yaghoobi and M.E. Davies, IEEE transactions on Aerospace and Electronic Systems, October 2013.

The phase and/or gain ambiguities in SAR imaging are modelled using sparse dictionary learning techniques to calibrate the parameters.

Abstract: Motivated by the field of compressed sensing and sparse recovery, nonlinear algorithms have been proposed for the reconstruction of synthetic aperture radar images when the phase history is under-sampled. These algorithms assume exact knowledge of the system acquisition model. In this paper we investigate the effects of acquisition model phase errors when the phase history is under-sampled. We show that the standard methods of autofocus, which are used as a postprocessing step on the reconstructed image, are typically not suitable. Instead of applying autofocus as a post-processor, we propose an algorithm that corrects phase errors during the image reconstruction. The performance of the algorithm is investigated quantitatively and qualitatively through numerical simulations on two practical scenarios where the phase histories contains phase errors and are under-sampled.

[P6] Sparsity-based Image Formation and RFI Mitigation for UWB SAR. S.I. Kelly, S.I. and M. Davies, NATO Specialist Meeting on Compressed Sensing for RADAR/SAR, Tallinn, Estonia, May 2014.

This work presented our recent advances in iterative reconstruction and RFI suppression for Low Frequency (UWB) Synthetic Aperture Radar.

[P7] A fast decimation-in-image back-projection algorithm for SAR, S. I. Kelly and M. E. Davies, IEEE Radar Conference, May 2014, pp. 1046-1051.

Motivated by the decimation-in-time FFT algorithms, we proposed a novel algorithm for SAR imaging to provide speed up compared with conventional back-projection algorithms. The proposed algorithm has advantages which relate to the way it manifests errors.
Abstract: Fast back-projection algorithms are required for new modalities of SAR, such as UWB SAR. In this paper we propose a novel algorithm which we call the fast decimation-in-image back-projection algorithm due to its relation to decimation-in-time FFT algorithms. It is the natural dual of existing fast back-projection algorithms which are related to decimation-in-frequency FFT algorithms. The proposed algorithm provides similar speed up to existing algorithms, however, it has additional advantages. The advantages relate to the way in which the algorithm manifests errors. The size and nature of the errors introduced in the proposed algorithm are more desirable than that of existing algorithms.

[P8] Parallel Processing of the Fast Decimation-in-image Back-projection Algorithm for SAR, Shaun Kelly, M.E. Davies, J.S. Thompson, In Proc. Sensor Signal Processing for Defence, September Edinburgh 2014.

In this paper, we investigated the use of multi-core processing and graphic processing units to speed up the proposed fast back-projection algorithm in [P6].
Abstract: Fast back-projection algorithms provide substantial speedup when compared with the standard back-projection algorithm. However in many potential near-field applications of synthetic aperture radar, further speedup is still required in order to make the application operationally feasible. In this paper we investigate the application of multi-core central processing units and graphic processing units, which are now standard on most scientific workstations, to further speed up a very recently proposed fast back-projection algorithm (the fast decimation-in-image back-projection algorithm).

[P9] A Sparse Regularized Model for Raman Spectral Analysis, D. Wu, M. Yaghoobi, S. Kelly M. E. Davies and R. Clewes, SSPD 2014, September 2014.

We proposed a sparsity-driven algorithm for Raman Spectroscopy in this work to help identify the components in the chemical mixtures and quantitatively analyse the concentrations as well.

Abstract: Raman spectroscopy has for a long time performed as a common analytical technique in spectroscopic applications. A Raman spectrum depends upon how efficiently a molecule scatters the incident light (electron rich molecules often produce strong signals) which results in difficulties for relating the spectrum to the absolute amounts of present substances. The spectrum is however a stable and accurate representation of the sample measured especially considering that each molecule is associated with a unique spectrum. State-of-the-art spectroscopic calibration methods include the principal component regression (PCR) and partial least squares regression (PLSR) methods which have been proved to be efficient regression methods to realise the quantitative analysis of Raman spectrum. In this paper we consider the problem of Raman spectra deconvolution to analyse the sample composition, as well as possible unknown substances. In particular, we propose a sparse regularized model as a complement to traditional regression methods by leveraging the components sparsity compared to the whole chemical library and the spectra sparsity, given that the chemical fingerprint of effectiveness of this sparse regularized model.

[P10] Fast Non-Negative Orthogonal Matching, Pursuit, M. Yaghoobi, D. Wu and M. E. Davies, IEEE Signal Processing Letters September 2015.

This work considers one particular class of sparse signals, i.e. the non-negative signals. The proposed method provides an accelerated method for recovering such non-negative representations.
WP1.

Abstract: One of the important classes of sparse signals is the non-negative signals. Many algorithms have already been proposed to recover such non-negative representations, where greedy and convex relaxed algorithms are among the most popular methods. The greedy techniques have been modified to incorporate the non-negativity of the representations. One of such modifications has been proposed for the Orthogonal Matching Pursuit (OMP), which first chooses positive coefficients and uses a non-negative optimisation technique as a replacement for the orthogonal projection onto the selected support. Beside the extra computational costs of the optimisation program, we do not benefit from the fast implementation techniques of OMP. These fast implementations are based on the matrix factorisations. We here first carefully investigate the cases we seeking a positive representation, using a pursuit algorithm. We will then propose a new implementation which truly incorporate the positivity constraint of the coefficients. As a result, we will also present a novel fast implementation of the Non-Negative OMP which is based on the QR decomposition and an iterative coefficients update. We will empirically show that such a modification can easily accelerate the implementation with a factor of ten in a reasonable size problem.

[P11] A Computationally Efficient Multi-coset Wideband Radar ESM Receiver, M Yaghoobi, M Davies, NATO Specialist Meeting on Compressed Sensing for RADAR/SAR, Tallinn, Estonia, May 2014.

Abstract: The problem of efficient sampling of wideband Radar signals for Electronic Support Measures (ESM) using a parallel sampling structure is investigated in this paper. Wideband radio frequency sampling, which is a necessary component of modern Radar surveillance systems, generally needs a sampling rate at least twice the maximum frequency of the signal, i.e. Nyquist rate, which is generally very high. However, when the signal is highly structured, like Radar signals, we can use the

fact that signals do not occupy the whole spectrum and instead, there exists a parsimonious structure in the time-frequency domain. Here, we characterise a novel low-complexity sampling system with a recovery guarantee, assuming that received RF signals have a particular structure. The proposed technique is inspired by the compressive sampling (CS) of sparse signals and it uses a multi-coset sampling setting, while it does not involve a computationally expensive reconstruction step. The new framework is comparable with current rapid frequency sweeping technique, while it continuously monitors the spectrum. It is thus affected from short pulse misidentification. On the other hand, like other CS based sub-Nyquist sampling techniques, it suffers from noise folding effect.

[P12] Compressed Sensing Solutions for Airborne Low Frequency SAR, M.E. Davies and Shaun Kelly, invited talk, EuRAD2014. Slides available at:
http://www.see.ed.ac.uk/drupal/sites/default/files/talk_MED.pdf

[P13] Implementation of an autocorrelation-based spectrum sensing algorithm in real-world channels with frequency offset, Pat Chambers, Mathini Sellathurai, Heriot-Watt University, SSPD2014, September 2014.

Abstract: This work presents a testbed implementation of a spectrum sensing algorithm for cognitive radio that is based on the autocorrelation function. Much of the work in current literature uses simulation based approaches to characterize functionality. In contrast here, the algorithm is applied in real-world channels and compared with appropriate simulations. It is shown how the algorithm may be improved to overcome the problem of frequency offset, which is a hardware-based impairment that current literature on the algorithm generally does not consider.

[P14] Passive Radar Signal Processing: Using OFDM based communications signals, M. Sellathurai, Military Radar, London, 2014 October 29th-30th.

Present automated extraction and classification of OFDM based reflected communications signals using computer vision and neural network based techniques.

Proposing a novel way of extracting the communications signals from a noisy spectrogram using a combination of morphological operations in conjunction with a fuzzy neighbourhood thresholding based bidirectional self-organizing neural network (BDSONN). Results showing about 98% detection is achieved at 5% false alarm at very low SNR outperforming traditional techniques.

[P15] "CFAR Analysis of the Multicoset-Thresholding Detector: Application to the Low Complexity Sub-Nyquist Radar Electronic Surveillance", M. Yaghoobi, B. Mulgrew, Andy Stove and Mike E. Davies, CoSeRa conference, June 2015.

Abstract: Multicoset sampling scheme is a technique to achieve high-speed sampling rate, using a bank of lower-rate sampling channels. In this technique, each channel samples with a small delay with respect to the other channels. As a result, we often can reconstruct the high bandwidth input signal, by wisely combining the information from different channels. However, in many applications, the reconstruction is not the goal. Here, we consider an application, i.e. Radar Electronic Surveillance, in which the aim is the detection and identification of the incoming Radar pulses. As the sampling rate is very high, e.g. up to tens of Giga samples per second, we also need a fast detection scheme. We have recently proposed an efficient multicoset sampling technique, called

LoCoMC, which is based on the thresholding for the detection and combining the information from different channels, to extract pulses. We here present an analytical investigation of the thresholding based detection and demonstrate how to choose the thresholding parameter. We then show that the algorithm can be competitive with a state of the art algorithm in performance, while it is computationally very cheap.

[P16] "Fast Non-Negative Orthogonal Least Squares", M. Yaghoobi and Mike E. Davies, EUSIPCO conference, August 2015.

Abstract: An important class of sparse signals is the non-negative sparse signals. While numerous greedy techniques have been introduced for low-complexity sparse approximations, there are few non-negative versions. Among such a large class of greedy techniques, one successful method, which is called the Orthogonal Least Squares (OLS) algorithm, is based on the maximum residual energy reduction at each iteration. However, the basic implementation of the OLS is computationally slow. The OLS algorithm has a fast implementation based on the QR matrix factorisation of the dictionary. The extension of such technique to the non-negative domain is possible. In this paper, we present a fast implementation of the non-negative OLS (NNOLS). The computational complexity of the algorithm is compared with the basic implementation, where the new method is faster with two orders of magnitude. We also show that, if the basic implementation of NNOLS is not computationally feasible for moderate size problems, the proposed method is tractable. We also show that the proposed algorithm is even faster than an approximate implementation of the non-negative OLS algorithm.

[P17] "Sparsity based Ground Moving Target Imaging via Multi-Channel SAR", Di Wu, Mehrdad Yaghoobi and Mike Davies, SSPD conference, September, 2015.

Abstract: State-of-the-art Ground Moving Target Indicator (GMTI) schemes include the Displaced Phase Center Antenna (DPCA) and Along Track Interferometry (ATI) which are commonly used image-based dual-channel techniques for moving target detection. In the present paper, we provide a different perspective for solving GMTI tasks by generalising the ground moving targets imaging as a parameter estimation and an optimisation problem. A sparsity based ground target imaging approach is described to improve the image quality for moving targets and estimate their states. By exploiting the fact that moving targets are highly sparse in the observed scene and feasible velocity space, the proposed method constructs a velocity map for the illuminated region, and combines this map with a sparsity based optimisation algorithm to realise the image formation. The performance of the presented method is demonstrated through GOTCHA airborne SAR data set.

[P18] Distributed fusion of PHD filters via Exponential Mixture Densities, M. Uney, D.E. Clark, S.J. Julier IEEE Journal of Selected Topics in Signal Processing, vol.7, no.3, pp.521—531, June 2013.

Abstract: In this paper, we consider the problem of Distributed Multi-sensor Multi-target Tracking (DMMT) for networked fusion systems. Many existing approaches for DMMT use multiple hypothesis tracking and track-to-track fusion. However, there are two difficulties with these approaches. First, the computational costs of these algorithms can scale factorially with the number of hypotheses. Second, consistent optimal fusion, which does not double count information, can only be guaranteed for highly constrained network architectures which largely undermine the benefits of distributed fusion. In this paper, we develop a consistent approach for DMMT by combining a generalised version of Covariance Intersection, based on Exponential Mixture Densities (EMDs), with Random Finite Sets (RFS). We first derive explicit formulae for the use of EMDs with RFSs. From this, we develop expressions for the probability hypothesis density filters. This approach supports DMMT in arbitrary network topologies through local communications and computations. We implement this approach using Sequential Monte Carlo techniques and demonstrate its performance in simulations.

[P19] A cooperative approach to sensor localisation in distributed fusion networks, *Murat Uney,* Bernard Mulgrew, Daniel Clark, IEEE Transactions on Signal Processing, accepted for publication.

Abstract: We consider self-localisation of networked sensor platforms which are located disparately and collect cluttered and noisy measurements from an unknown number of objects (or, targets). These nodes perform local filtering of their measurements and exchange posterior densities of object states over the network to improve upon their myopic performance. Sensor locations need to be known, however, in order to register the incoming information in a common coordinate frame for fusion. In this work, we are interested in scenarios in which these locations need to be estimated solely based on the multi-object scene. We propose a cooperative scheme which features nodes using only the information they already receive for distributed fusion: we first introduce node-wise separable parameter likelihoods for sensor pairs, which are recursively updated using the incoming multi-object information and the local measurements. Second, we establish a network coordinate system through a pairwise Markov random field model which has the introduced likelihoods as its edge potentials. The resulting algorithm consists of consecutive edge potential updates and Belief Propagation message passing operations. These potentials are capable of incorporating multi-object information without the need to find explicit object-measurement associations and updated in linear complexity with the number of measurements. We demonstrate the efficacy of our algorithm through simulations with multiple objects and complex measurement models.

[P20] Regional Variance for Multi-Object Filtering, Delande, E.; Uney, M.; Houssineau, J.; Clark, D IEEE Transactions on Signal Processing, vol.62, no.13, pp.3415,3428, July, 2014.

Abstract: Recent progress in multi-object filtering has led to algorithms that compute the first-order moment of multi-object distributions based on sensor measurements. The number of targets in arbitrarily selected regions can be estimated using the first-order moment. In this work, we introduce explicit formulae for the computation of the second-order statistic on the target number. The proposed concept of regional variance quantifies the level of confidence on target number estimates in arbitrary regions and facilitates information-based decisions. We provide algorithms for its computation for the probability hypothesis density (PHD) and the cardinalized probability hypothesis density (CPHD) filters. We demonstrate the behaviour of the regional statistics through simulation examples.

[P21] Cooperative sensor localisation in distributed fusion networks by exploiting non-cooperative targets, Murat Uney, Bernard Mulgrew, Daniel Clark, IEEE Statistical Signal Processing Workshop, June 2014, 29/06-02/07, Gold Coast, Australia.

Abstract: We consider geographically dispersed and networked sensors collecting measurements from multiple targets in a surveillance region. Each sensor node filters the set of cluttered, noisy target

measurements it collects in a sensor centric coordinate system and with imperfect detection rates. The filtered multi-target information is, then, communicated to the nearest neighbours. We are interested in network self-localisation in scenarios in which the network is restricted to use only the multi-target information shared. We propose an online distributed sensor localisation scheme based on a pairwise Markov Random Field model of the problem. We first introduce parameter likelihoods for pairs of sensors -equivalently, edge potentials- which can be computed using only the incoming multi-target information and local measurements. Then, we use belief propagation with the associated posterior model which is Markov with respect to the underlying communication topology. We demonstrate the efficacy of our algorithm for cooperative sensor localisation through an example with complex measurement models.

[P22] Target aided online sensor localisation for bearing only clusters, Murat Uney, Bernard Mulgrew, Daniel Clark, Sensor Signal Processing in Defence Conference, September 2014, Edinburgh, UK.

Abstract: In this work, we consider a network of bearing only sensors in a surveillance scenario. The processing of target measurements follow a two-tier architecture: The first tier is composed of centralised processing clusters whereas in the second tier, cluster heads perform decentralised processing. We are interested in the first tier problem of locating peripheral sensors relative to their cluster head. We mainly exploit target measurements received by the cluster head in a parameter estimation setting which involves Sequential Monte Carlo methods, and is known to have many difficulties in practice, including particle deficiency, sensitivity to initialisation, and high computational complexity. These difficulties are exacerbated by the bearing-only modality which provides a relatively poor target observability. We propose an online solution through Bayesian recursions on Junction Tree models of the posterior which partition the problem into fixed size subproblems and hence provides scalability with the number of sensors. We use the received signal strength as noisy range measurements to improve the robustness and accuracy of our algorithm. We demonstrate its efficacy with an example.

[P23] Regional variance in target number: Analysis and application for multi-Bernoulli point processes, Delande, E.D.; Houssineau, J.; Clark, D.E., Data Fusion & Target Tracking 2014: Algorithms and Applications (DF&TT 2014), IET Conference, pp.1,8, 30-30 April 2014.

Abstract: In the context of multi-target tracking application, the concept of variance in the number of targets estimated in specified regions of the surveillance scene has been recently introduced for multi-object filters. This article has two main objectives. First, the regional variance is derived for a multi-object representation commonly used in the tracking literature, known as the multi-Bernoulli point process, in which the multi-target state is described with a set of hypothesised tracks with associated existence probabilities. This model is exploited in multi-target applications where it can be assumed that targets evolve independently of each other and generate sensor observations that are uncorrelated with other targets. An illustration of the concept of regional statistics (mean and variance) in target number, and how to interpret them in the broader context of multi-object filtering, it then provided. Possible applications include performance assessment and sensor control for multi-target tracking.

[P24] M. Uney, B. Mulgrew, D.E. Clark, "Maximum likelihood signal parameter estimation via track before detect," in Sensor Signal Processing for Defence (SSPD) September 2015, Edinburgh, UK.

Abstract: In this work, we consider the front-end processing for an active sensor. We are interested in estimating signal amplitude and noise power based on the outputs from filters that match transmitted waveforms at different ranges and bearing angles. These parameters identify the distributions in, for example, likelihood ratio tests used by detection algorithms and characterise the probability of detection and false alarm rates. Because they are observed through measurements induced by a (hidden) target process, the associated parameter likelihood has a time recursive structure which involves estimation of the target state based on the filter outputs. We use a track-before-detect scheme for maintaining a Bernoulli target model and updating the parameter likelihood. We use a maximum likelihood strategy and demonstrate the efficacy of the proposed approach with an example.

[P25] M. Uney, B. Mulgrew, D.E. Clark, "Distributed Estimation of Latent Parameters in State Space Models Using Separable Likelihoods," submitted to ICASSP 2016.

Abstract: Motivated by object tracking applications with networked sensors, we consider multi sensor state space models. These models are associated with latent parameters which must be estimated using noisy sensor measurements. Because the measurements are induced by (hidden) Markov processes, the parameter likelihoods depend on measurement histories of more than one sensor and require centralisation. We propose an approximation with a node-wise separable structure thereby removing the need for centralisation in computing likelihoods. We establish the connection between quality of approximation and the accuracy of state estimation based on individual sensor histories. When leveraged with Markov random field models and message passing algorithms for inference, these likelihoods facilitate decentralised estimation in tracking networks as well as a scalable computation schemes in centralised settings. We demonstrate this approach in a sensor network self-localisation example.

[P26] Independent views in MIMO sonar systems, Yan Pailhas and Yvan Petillot, Underwater Acoustics, June 2014.

Abstract: The main advantages of MIMO (Multiple Input Multiple Output) sonar systems come from the assumption of independent observations between each transmitter/receiver pairs. The independence of the observations ensures an unbiased set of measurements and then provides true statistics on the target. In this paper we study the correlation between views in MIMO sonar systems. A traditional tool used to study the dependency between two random variables is the Pearson product-moment correlation coefficient. However this measure suffers numerous defaults: it only estimates linear correlation, it is not a proper distance and in particular a null measure of the Pearson coefficient does not insure the independence of the tested random variables. For these reasons we will use the distance correlation introduced by Szekely. From the distance correlation we will derive the inter-views distance correlation matrix which assess the correlation of the full MIMO system (i.e. the dependencies between each views). This independence measure matrix gives a guideline to how to build truly uncorrelated MIMO sonar systems and then maximise the performances of such system.

[P27] Tracking underwater objects using large MIMO sonar systems, Yan Pailhas, Jeremie Houssineau, Emmanuel Delande, Jeremie Houssineau, Yvan Petillot and Daniel Clark, Underwater Acoustics, June 2014.

Abstract: MIMO sonar systems can offer great capabilities for area surveillance especially in very shallow water with heavy cluttered environment. We present here a MIMO simulator which can compute synthetic raw data for any transmitter/receiver pair in multipath and cluttered environment. Synthetic moving targets such as boats or AUVs can also be introduced into the environment. For the harbour surveillance problem we are interested in tracking all moving objects in a particular area. So far the tracking filter of choice for multistatic systems has been the MHT (Multiple Hypothesis Tracker). The reason behind this choice is its capability to propagate track identities at each iteration. The MHT is an extension of a mono object tracker to a multi object problem and therefore suffers from a number of drawbacks: the number of targets should be known and the birth or death of new tracks are based on heuristics. A fine ad hoc parameter tuning is then required and there is a lack of adaptivity in this process. To overcome those restrictions we will be using the HISP (Hypothesised multi-object filter for Independent Stochastic Population) filter recently developed. The HISP filter relies on a generalisation of the concept of point process that integrates a representation of distinguishability. As a consequence, this filter deals directly with the multi-object estimation problem, while maintaining track identities through time without using heuristics. While filters track the objects after processing in the digital domain, we show as well in this paper that we can adapt acoustical time reversal techniques to track an underwater target directly with the MIMO system. We will show that the proposed modified DORT technique matches the prediction / data update steps of a tracking filter.

[P28] Large MIMO sonar systems: a tool for underwater surveillance, Yan Pailhas, Yvan Petillot, SSPD September 2014.

Abstract: Multiple Input Multiple Output sonar systems offer new perspectives for target detection and underwater surveillance. In this paper we present an unified formulation for sonar MIMO systems and study their properties in terms of target recognition and imaging. Here we are interested in large MIMO systems. The multiplication of the number of transmitters and receivers non only provides a greater variety in term of target view angles but provides also in a single shot meaningful statistics on the target itself. We demonstrate that using large MIMO sonar systems and with a single shot it is possible to perform automatic target recognition and also to achieve super-resolution imaging. Assuming the view independence between the MIMO pairs the speckle can be solved and individual scatterers within one resolution cell decorelate. A realistic 3D MIMO sonar simulator is also presented. The output of this simulator will demonstrate the theoretical results.

[P29] Synthetic aperture imaging and autofocus with coherent MIMO sonar systems, Yan Pailhas, Yvan Petillot, IOA SAR/SAS conference September 2014.

Abstract: MIMO stands for Multiple Inputs Multiple Outputs. Such systems have received a lot of interests in the radar community during the last decade. One of the main reason behind this is the greater variety in terms of target view angles compared with traditional monostatic systems. With several independent views one can hope to reduce the speckle effect typical of coherent sensor systems such as RADAR or SONAR. Destructive interference in particular can be the cause of missed detection. However MIMO sonar systems have been studied by the ASW (anti-submarine warfare) community mainly to increase the probability of detection of low target strength target. Few MIMO

systems have been built and tested including DEMUS (LF SIMO system built by CMRE). In any case all those systems use a relatively low number of sensors. With large MIMO systems we proved that it is possible to solve the speckle within one resolution cell and then archive super-resolution images making such systems very attractive for surveillance. The independent view assumption places constraints on the MIMO design especially on the sensor locations. The sensor locations for MIMO systems are then extremely sparse compared to a lambda/2 phased array and the sidelobes can be significant. The main tool for MIMO imaging is based on the back projection algorithm extended to multi static systems. We discuss sidelobe reduction for MIMO imaging using randomised sensor positioning and high peak suppression. We also demonstrate that autofocus techniques can be applied to estimate with great accuracy mid water target depth and speed. All the MIMO data in this paper are computed using a full 3D realistic MIMO simulator including multipath, seabed physical models and cloud point model to compute time echoes.

[P30] MIMO Sonar Systems for Harbour Surveillance, Yan Pailhas, Yvan Petillot, Proceeding of the IEEE Oceans Conference, October 2015.

Abstract: The MIMO acronym stands for Multiple Input Multiple Output. It refers to a system with several transmitters and several receivers. MIMO systems can been seen as a variety of multi-static systems, the main difference being that MIMO system has the capability to process the information as a whole while multi-static systems only process the data at the receiver nodes. This implies that there is an overall strategy for MIMO systems, a strategy specific to the end application. We can distinguish two levels of freedom in MIMO systems, the first one being at the transmitter level. The classic approach is to consider orthogonal waveforms in order to separate the bistatic signals from all the different transmitters. Finite orthogonal waveforms do not exist and the search for approximate orthogonal waveforms is an active subject of research. The other level of freedom is at the receiver end and different techniques can be designed to extract target information from the N x M signals (where N is the number of transmitters and M the number of receivers). Assuming a coherent system we will demonstrate that with a high number of independent observations MIMO sonar systems solve the speckle and can achieve super resolution imaging. In this paper we are interested in the problem of harbour surveillance. MIMO sonar systems can offer great capabilities for area surveillance especially in very shallow water with heavy cluttered environment such as harbour environment. To benefit from the view diversity we consider here MIMO systems with spatially distributed transmitters and receivers. A full 3D MIMO simulator will be presented which can compute synthetic raw data for any transmitter/receiver pair in multipath and cluttered environment. Synthetic seabed interface are computed using 2D fractional Brownian motion. Bistatic reverberation level are computed using physical model developed by APL-UW. Finally mirror theorem is used to compute the various multipaths. Synthetic mid-water targets can also be added to the environment. However the sound propagation in 3D can be computationally expensive. We present here sparse techniques which reduce the computation time drastically. Assuming a coherent MIMO sonar system, MIMO image formation can be processed using multi-static back projection algorithm (variant of the bistatic back projection algorithm developed by the SAR community). For the multi-static scenario the continuous integration along ellipses is replaced by a finite sum in which each term corresponds to one transmitter/receiver pair contribution. In this paper we propose three variants of the multistatic back projection algorithm with autofocus capability to calibrate the MIMO array, to estimate the depth of a mid-water target and to estimate its speed and orientation. Finally we will show that large MIMO systems offer an ideal platform for time reversal techniques if we relax the orthogonal waveform

assumption. We will present in particular an unfocussed time reversal mirror algorithm capable of automatically tracking moving targets. Despite the physical constraints, the high number of transmitters and receivers give the user a great degree of freedom on how to use and exploit large MIMO sonar systems. In this paper we present a series of techniques based on autofocus and defocus which allow a MIMO sonar system in a harbour environment to automatically track or estimate number of parameters such as speed or depth from a mid water target.

[P31] Multi dimensional Fast Marching Approach to Wave Propagation in Heterogenous Multipath Environment, Yan Pailhas, Yvan Petillot, UACE15, June 2015.

Abstract: Sound propagation is described by the wave equation. If in an homogenous free field its resolution is straightforward, any variation from this hypothesis makes the wave equation solution not analytically tractable especially in a shallow water environment. The direct resolution of the wave equation requires in most cases numerical methods such as FDTD (Finite Difference Time Domain) or PSTD (Pseudo Spectral Time Domain). A direct approach however is often extremely computationally expensive (the spatial and temporal discretisation has to be small, of the order of $\lambda/10$, for stability criterion) and approximations are necessary for practical reasons. For low frequencies applications, mathematical models for shallow water propagation include Normal Mode Model or Parabolic Equation Model. For higher frequencies (above 1 kHz), the most popular method for wave propagation in shallow water is based on Ray theory and geometrical acoustics. Thanks to the infinite frequency assumption, the wave equation simplifies to the eikonal equation which propagates the wavefront of the acoustic pulse. The ray trajectories are computed as perpendicular to the wavefront. In the ideal case of a constant sound velocity profile and perfectly flat interfaces for the surface and the seafloor, an elegant solution is derived from the Mirror theorem: source images are easily geometrically computed by successive symmetries of the source itself. In a second step folding the straight paths linking all the source images to a target computes the multipath. Unfortunately this method fails for non flat seabeds, non constant depth or non constant velocity profile. In this paper we propose an extension to the Mirror theorem to take into account any interface geometry or sound velocity variation (horizontally or vertically) by solving the eikonal equation using the Fast Marching algorithm. We will show that multipath can then be solved by wrapping the wavefront propagation at each interface.

[P32] Review on Orthogonal Waveforms for Large MIMO Sonar Applications, Yan Pailhas, Yvan Petillot, UACE15, June 2015.

Abstract: MIMO means Multiple Inputs Multiple Outputs. Such systems have been developed at first for radar applications. MIMO recently have gained interest in the underwater acoustic community because of certain benefits over traditional systems such as increase resolution or increase in signal to clutter ratio to name a few. The MIMO concept relies on multiple transmitters (Nt) sending unique and orthogonal waveforms through the environment. Several receivers (Nr) then capture environment, targets or clutter echoes. At each receiver point the total signal is filtered to separate each transmitter signal. The stage at which MIMO systems separates from multi-static systems is the information processing, which is done centrally rather that separately at each receiver node. Accessing the Nt x Nr signals requires the orthogonality of the out-coming pulses. As purely orthogonal waveforms do not exist, different approaches were developed to minimise the waveform cross-correlation. Such methods include CDMA (code division multiple access) where waveforms share the

same frequencies at the same time, TDMA (time division multiple access) where waveforms share the same frequency band, but at different times, or FDMA (frequency division multiple access) where waveforms occupy different frequencies at the same time. In this paper we review the three main classes of orthogonal waveforms and present preliminary results in a test tank and real environment.

[P33] Spatially distributed MIMO sonar systems: principles and capabilities, Yan Pailhas, Yvan Petillot, Keith Brown, Bernie Mulgrew, IEEE Journal of Oceanic Engineering (submitted).

Abstract: Multiple Input Multiple Output sonar systems offer new perspectives for target detection and area surveillance. This paper introduces an unified formulation for sonar MIMO systems and focuses on the target detection and recognition capability of these systems. The multiplication of the number of transmitters and receivers not only provides a greater variety in terms of target view angles but provides also in meaningful statistics on the target itself. Assuming that views are independent and the MIMO system is large enough we demonstrate that target recognition is possible with only one view from the full system. By studying the detection performance of MIMO sonars we also demonstrate that such systems solve the speckle noise and decorrelate individual scatterers inside one cell resolution. We show that MIMO systems can achieve super-resolution images and surpass the resolution given by equivalent SAS (Synthetic Aperture Sonar) systems. All the discussed properties are derived from the independent view assumption. This assumption is discussed and leads to the design requirement and efficiency of MIMO sonar systems.

[P34] Wideband CDMA waveforms for large MIMO sonar systems, Yan Pailhas and Yvan Petillot, SSPD conference 2015, Edinburgh, UK, September, 2015.

Abstract: Multiple Input Multiple Output (MIMO) sonar systems offer new perspectives for target detection and underwater surveillance. The inherent principle of MIMO relies on transmitting several pulses from different transmitters. The MIMO waveform strategy can vary from applications to applications. But among the waveform space, orthogonal waveforms are arguably the most important sub-space. Purely orthogonal waveforms do not exist, and several approximations have been attempted for MIMO radar applications. These approaches include separating the waveforms in the time domain, the frequency domain or using pseudo orthogonal codes. In this paper we discuss the different radar waveform approaches from a sonar point of view and propose a novel CDMA (code division multiple access) waveform design, more suitable for large wideband MIMO systems.

[P35] Accelerating the Single Cluster PHD Filter with a GPU Implementation, Chee Sing Lee, Jose Franco, Jérémie Houssineau, Daniel Clark, International Conference on Control, Automation and Information Sciences (ICCAIS), December 2014

Abstract: The SC-PHD filter is an algorithm which was designed to solve a class of multiple object estimation problems where it is necessary to estimate the state of a single-target parent process, in addition to estimating the state of a multi object population which is conditioned on it. The filtering process usually employs a number of particles to represent the parent process, coupled each with a conditional PHD filter, which is computationally burdensome. In this article, an implementation is described which exploits the parallel nature of the filter to obtain considerable speed-up with the help of a GPU. Several considerations need to be taken into account to make efficient use of the GPU, and these are also described here.

[P36] Anomaly detection in clutter using spectrally enhanced Ladar, Puneet S. Chhabra, Andrew M. Wallace and James R. Hopgood (accepted SPIE conference, April 2015).

Abstract: Discrete return Laser Detection and Ranging (Ladar) systems provide a series of echoes (first/last or multiecho) that reflect from objects in a scene. On the other hand, Full-Waveform (FW)-Ladar systems measure the intensity of light reflected from objects continuously over a period of time and has not been often applied to anomaly detection. This paper presents an anomaly detection algorithm that considers Multi-Spectral (MS) FWLadar (MSL) measurements as a set of multi-dimensional data samples. We present a framework that allows the detection of spectral and time anomalies in FW-MSL data. In the signal domain we define an anomaly as a full waveform time and spectral signature that does not conform to a prior expectation, defined using a learnt subspace (dictionary) and co-occurring local-patterns. We propose an optimization algorithm for subspace learning based on stochastic approximations and augment our objective function with a discriminative term that represents the subspace's separability properties.

[P37] Information Processing for Foliage Penetrating LiDAR, Chhabra, P., ICVSS, Sicily, Italy, 13-19th July 2014.

Abstract: Discrete return LiDAR systems provide a series of echoes (first/last or multi-echo) that reflect targets in a scene. On the other hand, a Full-Waveform (FW) LiDAR system measures the intensity of light that reflects targets continuously over a period of time. Research relating to FW-LiDAR is fairly new and barely scratched for target detection, surveillance and combat identification using Multi-Spectral (MS) FW-LiDARs. This ongoing work addresses the following problem: How best to combine and filter point cloud data acquired from ground based/aerial spectrally enhanced FW-LiDAR sensors to create detailed situational awareness.

[P38] Human behaviour recognition in data-scarce domains. RH Baxter, NM Robertson, DM Lane, Pattern Recognition, August 2015.

Abstract: This paper presents the novel theory for performing multi-agent activity recognition without requiring large training corpora. The reduced need for data means that robust probabilistic recognition can be performed within domains where annotated datasets are traditionally unavailable. Complex human activities are composed from sequences of underlying primitive activities. We do not assume that the exact temporal ordering of primitives is necessary, so can represent complex activity using an unordered bag. Our three-tier architecture comprises low-level video tracking, event analysis and high-level inference. High-level inference is performed using a new, cascading extension of the Rao-Blackwellised Particle Filter. Simulated annealing is used to identify pairs of agents involved in multi-agent activity. We validate our framework using the benchmarked PETS 2006 video surveillance dataset and our own sequences, and achieve a mean recognition F-Score of 0.82. Our approach achieves a mean improvement of 17\% over a Hidden Markov Model baseline.

[P39] An adaptive motion model for person tracking with instantaneous head-pose features. RH Baxter, MJV Leach, SS Mukherjee, NM Robertson, IEEE Signal Processing Letters, 22(5), pp 578-582. February 2015.

Abstract: This paper presents novel behaviour-based tracking of people in low-resolution using instantaneous priors mediated by head-pose. We extend the Kalman Filter to adaptively combine

motion information with an instantaneous prior belief about where the person will go based on where they are currently looking. We apply this new method to pedestrian surveillance, using automatically-derived head pose estimates, although the theory is not limited to head-pose priors. We perform a statistical analysis of pedestrian gazing behaviour and demonstrate tracking performance on a set of simulated and real pedestrian observations. We show that by using instantaneous `intentional' priors our algorithm significantly outperforms a standard Kalman Filter on comprehensive test data.

[P40] Robust joint audio-video tracking. E D'Arca, NM Robertson, JR Hopgood. ACM Transactions on Multimedia *(submitted 2015).*

Abstract: Situational awareness is achieved naturally by the human senses of sight and hearing in combination. System level automatic scene understanding aims at replicating this human ability using cooperative microphones and cameras. In this paper, we integrate and fuse audio and video signals at different levels of abstractions to detect and track a speaker in a scenario where people are free to move indoors. Despite the low complexity of the system, which consists of just 4 microphone pairs and 1 camera, results show that the overall multimodal tracker is more reliable than single modality systems, tolerating large occlusions and cross-talking. The system evaluation is performed on both single modality and multimodality tracking. The performance improvement given by the audio-video integration and fusion, is quantified in terms of tracking precision and accuracy as well as speaker diarisation error rate and precision-recall recognition metrics. We evaluate improvements vs. the closest works: 56% on audio only sound source localisation computational cost, 18% on the speaker diarisation error rate over an audio only speaker recognition unit and 36% on the precision-recall metric over an audio-video dominant speaker recognition method.

[P41] Tracking with intent. RH Baxter, M Leach, NM Robertson. Sensor Signal Processing for Defence (SSPD), September 2014.

Abstract: This paper presents the novel theory for performing behaviour-based tracking using intentional priors. Motivated by our ultimate goal of anomaly detection, our approach is rooted in building better models of target behaviour. Our novel extension of the Kalman filter combines motion information with an intentional prior. We apply our `Intentional Tracker' to a pedestrian surveillance and tracking problem, using head pose as the intentional prior. We perform a statistical analysis of pedestrian head pose behaviour and demonstrate tracking performance on a set of simulated and real pedestrian observations. We show that by using intentional priors our algorithm outperform a standard Kalman filter across a range of target trajectories.

[P42] Detecting social groups in crowded surveillance videos using visual attention. Michael Leach, Rolf Baxter, Neil Robertson, Ed Sparks, Computer Vision and Pattern Recognition Workshops, pp 467-473, June 2014.

Abstract: In this paper we demonstrate that the current state of the art social grouping methodology can be enhanced with the use of visual attention estimation. In a surveillance environment it is possible to extract the gazing direction of pedestrians, a feature which can be used to improve social grouping estimation. We implement a state of the art motion based social grouping technique to get a baseline success at social grouping, and implement the same grouping with the addition of the visual attention feature. By a comparison of the success at finding social groups for two techniques we evaluate the effectiveness of including the visual attention feature. We test both methods on two datasets containing busy surveillance scenes. We find that the inclusion of visual interest improves the motion social grouping capability. For the Oxford data, we see a 5.6% improvement in true positives and 28.5% reduction in false positives. We see up to a 50% reduction in false positives in other datasets. The strength of the visual feature is demonstrated by the association of social connections that are otherwise missed by the motion only social grouping technique.

[P43] Look who's talking: Detecting the dominant speaker in a cluttered scenario. E D'Arca, NM Robertson, JR Hopgood. IEEE International Conference on Acoustics, Speech and Signal Processing, May 2014.

Abstract: In this work we propose a novel method to automatically detect and localise the dominant speaker in an enclosed scenario by means of audio and video cues. The underpinning idea is that gesturing means speaking, so observing motions means observing an audio signal. To the best of our knowledge state-of-the-art algorithms are focussed on stationary motion scenarios and close-up scenes where only one audio source exists, whereas we enlarge the extent of the method to larger field of views and cluttered scenarios including multiple non-stationary moving speakers. In such contexts, moving objects which are not correlated to the dominant audio may exist and their motion may incorrectly drive the audio-video (AV) correlation estimation. This suggests extra localisation data may be fused at decision level to avoid detecting false positives. In this work, we learn Mel-frequency cepstral coefficients (MFCC) coefficients and correlate them to the optical flow. We also exploit the audio and video signals to estimate the position of the actual speaker, narrowing down the visual space of search, hence reducing the probability of incurring in a wrong voice-to-pixel region association. We compare our work with a state-of-the-art existing algorithm and show on real datasets a 36% precision improvement in localising a moving dominant speaker through occlusions and speech interferences.

[P44] Dynamic Distance-based Shape Features for Gait Recognition, Journal of Mathematical Imaging and Vision, March 2014, T.Whytock, A.Belyaev, N.M.Robertson.

Abstract: We propose a novel skeleton-based approach to gait recognition using our Skeleton Variance Image. The core of our approach consists of employing the screened Poisson equation to construct a family of smooth distance functions associated with a given shape. The screened Poisson distance function approximation nicely absorbs and is relatively stable to shape boundary perturbations which allows us to define a rough shape skeleton. We demonstrate how our Skeleton Variance Image is a powerful gait cycle descriptor leading to a significant improvement over the existing state of the art gait recognition rate.

[P45] On covariate factor detection and removal for robust gait recognition, Machine Vision and Applications, March 2015, T.Whytock, A.Belyaev, N.M.Robertson (accepted).

Abstract: Robust gait recognition is imperative to overcome covariate factors such as clothing, bags, shoes and elapsed time between capture. Our approach detects covariate factors and we develop and evaluate three techniques, varying in aggression, to remove their influence from both training and test data; this ensures their visual similarity and boosts performance by ensuring classification is based only on covariate factor free areas. Validation is performed on two of the largest and covariate factor rich databases where we yield new state of the art results.

[P46] Video Tracking through Occlusions by fast audio source localisation." E. D'Arca, A. Hughes, N. M. Robertson, J. Hopgood. IEEE Int. Conf. on Image Processing, Melbourne, September 2013.

Abstract: In this paper we present a novel audio-visual speaker detection and localisation algorithm. Audio source position estimates are computed by a novel stochastic region contraction (SRC) audio search algorithm for accurate speaker localisation. This audio search algorithm is aided by available video information (stochastic region contraction with height estimation (SRC-HE)) which estimates head heights over the whole scene and gives a speed improvement of 56% over SRC. We finally combine audio and video data in a Kalman filter (KF) which fuses person-position likelihoods and tracks the speaker. Our system is composed of a single video camera and 16 microphones. We validate the approach on the problem of video occlusion i.e. two people having a conversation have to be detected and localised at a distance (as in surveillance scenarios vs. enclosed meeting rooms). We show video occlusion can be resolved and speakers can be correctly detected/ localised in real data. Moreover, SRC-HE based joint audio-video (AV) speaker tracking outperforms the one based on the original SRC by 16% and 4% in terms of multi object tracking precision (MOTP) and multi object tracking accuracy (MOTA). Speaker change detection improves by 11% over SRC.

[P47] Using the voice spectrum for improved tracking of people in a joint audio-video scheme, E.D'Arca, N.M. Robertson and J. Hopgood, IEEE Int. Conf. Acoustics Speech and Signal Processing (ICASSP), Vancouver, May 2013.

Abstract: In this paper we present a new solution to the problem of speaker tracking among people where occlusions occur (disappearance and non-speaking). In a normal conversation between two or more people, we learn speaker mel-cepstral coefficients (MFCC) and incorporate this information into a sequential Bayesian audio-video position tracker. The joint video-to-audio data association step is thus improved and we achieve robust person recognition which in turn aids tracking performance. We provide comprehensive evaluation via simulations and real data quoting tracking accuracy, precision and diarisation error rate (DER) compared to ground truth. For simulate and real experiments in an open space the trajectory tracking performance increases by 20% measured against ground truth using our approach. As a further enhancement versus the state-of-the-art, speaker identity recognition at a distance is improved by 20% by exploiting audio-video localisation cues.

[P48] Contextual Anomaly Detection in Crowded Surveillance Scenarios, M.Leach, E.Sparks and N.M.Robertson, Pattern Recognition Letters, December 2013 (doi: 0.1016/j.patrec.2013.11.018)

Abstract: This work addresses the problem of detecting human behavioural anomalies in crowded surveillance environments. We focus in particular on the problem of detecting subtle anomalies in a behaviourally heterogeneous surveillance scene. To reach this goal we implement a novel unsupervised context-aware process. We propose and evaluate a method of utilising social context and scene context to improve behaviour analysis. We find that in a crowded scene the application of Mutual Information based social context permits the ability to prevent self-justifying groups and propagate anomalies in a social network, granting a greater anomaly detection capability. Scene context uniformly improves the detection of anomalies in both datasets. The strength of our

contextual features is demonstrated by the detection of subtly abnormal behaviours, which otherwise remain indistinguishable from normal behaviour.

[P49] Look Who's Talking, E.D'Arca, N.M.Robertson, J.R.Hopgood, IET Conf. Intelligent Signal Processing, London, December 2013.

Abstract: This paper proposes a method to automatically detect and localise the dominant speaker in a conversation by using audio and video information. The idea is that gesturing means speaking, so we look for people hands or heads movements to infer a person is talking. In a normal conversational context with two or more people, we learn Mel-frequency cepstral coefficients (MFCC) and find how they correlate with the optical flow associated with moving pixel regions by
canonical correlation analysis (CCA). In complex scenarios, this operation could be resulting in associating pixel regions to sounds which actually are not really correlated. Therefore, we also triangulate the information coming from the microphones to estimate the position of the actual audio source, narrowing down the visual space of search, hence reducing the probabilities of incurring in a wrong voice-to-pixel region association. We compare our work with a state-of-the-art existing algorithm and show on real data the improvement in dominant speaker localization.
WP5

[P50] Instantaneous real-time head pose at a distance, S. S. Mukherjee, R. H. Baxter and Neil M. Robertson. IEEE International conference on image processing, Quebec, September 2015.

In this paper we focus on robust, real-time human head pose estimation in low resolution RGB data without any smoothing motion priors e.g. direction of motion. Our main contributions lie in three major areas. First, we show that a generative Deep Belief Network model can be learned on human head data from multiple types of data sources. These sources have similar underlying data that are not necessarily labelled or have the same kind of ground truth. Second, we perform discriminative training using multiple disparate supervisory labels to fine tune the model for head pose estimation. Third, we present state-of-the-art results on two publicly available datasets using this new approach. Our implemetation computes head pose for a head image in 0.8 milliseconds, making it real-time and highly scalable.

[P51] Deep Head Pose: gaze-direction estimation in multimodal video, S. S. Mukherjee and N. M. Robertson, IEEE Trans. Multimedia, Special issue on Deep Learning for Multimedia Computing, 2015.

In this paper we present a Convolutional Neural Network based model for human head pose estimation in lowresolution multi-modal RGB-D data. We pose the problem as one of classification of human gazing direction. We further fine-tune a regressor based on the learned deep classifier. Next we combine the two models (classification and regression) to estimate approximate regression confidence. We present state-ofthe-art results in datasets that span the range of high resolution Human Robot Interaction (close up faces plus depth information) data to challenging low resolution outdoor surveillance data. We build upon our robust head-pose estimation and further introduce a new visual attention model to recover interaction with the environment. Using this probabilistic model we show that many higher level scene understanding like human-human/scene interaction detection can be achieved. Our solution runs in realtime on commercial hardware.

[P52] A filter for distinguishable and independent populations, Emmanuel Delande, Jeremie Houssineau, Daniel E. Clark, ArXiV, 19 Jan 2015.

Abstract: This article introduces a multi-object filter for the resolution of joint detection/tracking problems involving multiple targets, derived from the novel Bayesian estimation framework for stochastic populations. Fully probabilistic in nature, the filter for Distinguishable and Independent Stochastic Populations (DISP) exploits two exclusive probabilistic representations for the potential targets. The distinguishable targets are those for which individual information is available through past detections; they are represented by individual tracks. The indistinguishable targets are those for which no individual information is available yet; they are represented collectively by a single stochastic population. Assuming that targets are independent, and adopting the "at most one measurement per scan per target" rule, the DISP filter propagates the set of all possible tracks, with associated credibility, based on the sequence of measurement sets collected by the sensor so far. A few filtering approximations, aiming at curtailing the computational cost of a practical implementation, are also discussed.

[P53] Performance metric in closed-loop sensor management for stochastic populations, Delande, E.D.; Houssineau, J.; Clark, D.E., Sensor Signal Processing for Defence (SSPD), 2014 , vol., no., pp.1,5, 8-9 Sept. 2014.

Abstract: Methods for sensor control are crucial for modern surveillance and sensing systems to enable efficient allocation and prioritisation of resources. The framework of partially observed Markov decision processes enables decisions to be made based on data received by the sensors within an information-theoretic context. This work addresses the problem of closed-loop sensor management in a multi-target surveillance context where each target is assumed to move independently of other targets. Analytic expressions of the information gain are obtained, for a class of exact multi-target tracking filters are obtained and based on the Rényi divergence. The proposed method is sufficiently general to address a broad range of sensor management problems through the application-specific reward function defined by the operator.

[P54] Localised variance in target number for the Cardinalized Probability Hypothesis Density filter, Emmanuel Delande, Jeremie Houssineau and Daniel Clark. International Conference on Information Fusion, July 2013.

Abstract: Following a recent study on the Probability Hypothesis Density filter, this paper aims at extracting higher-order information statistics on the local target number from the filtered state of the Cardinalized Probability Hypothesis Density filter, based on recent developments of novel derivation tools in the multi-object filtering framework. In addition to the description of a novel approach for retrieving the expression of the updated localised mean target number, this paper proposes the extraction of the novel localised variance in the target number across the whole state space.

[P55] PHD filtering with localised target number variance, Emmanuel D Delande, Jeremie Houssineau, Daniel E Clark, Signal Processing, Sensor Fusion, and Target Recognition XXII, April 2013.

Abstract: Mahler's Probability Hypothesis Density (PHD filter), proposed in 2000, addresses the challenges of the multiple-target detection and tracking problem by propagating a mean density of the targets in any region of the state space. However, when retrieving some local evidence on the target presence becomes a critical component of a larger process - e.g. for sensor management purposes - the local target number is insufficient unless some confidence on the estimation of the number of targets can be provided as well. In this paper, we propose a first implementation of a PHD filter that also includes an estimation of localised variance in the target number following each update step; we then illustrate the advantage of the PHD filter + variance on simulated data from a multiple-target scenario.

[P56] A Forward-Backward Cardinalized Probability Hypothesis Density Smoother, Daniel Clark, Sharad Nagappa, Emmanuel Delande, and Jeremie Houssineau, IEEE Transactions on Aerospace and Electronic Systems (Submitted 4/12/2014).

Abstract: Multi-object smoothing provides a means of reducing the error in multi-object filtering. The PHD smoother provides the first-order approximation to the multi-target Bayes smoother. The PHD smoother does not always improve the cardinality estimate, thus motivating the derivation of a cardinalized PHD (CPHD) smoother. We derive a tractable form of the CPHD smoother without target birth in this paper and show that it addresses the shortcomings of the PHD smoother in the cardinality estimation.

[P57] Faà Di Bruno's formula and volterra series, Clark, D.E.; Houssineau, J., Statistical Signal Processing (SSP), 2014 IEEE Workshop on , vol., no., pp.217,219, June 29 2014-July 2 2014.

Abstract: Volterra series are used for modelling nonlinear systems with memory effects. The nth-order impulse response and the kernels in the series can be determined with Fréchet derivatives of Volterra series operators. Consequently, we can determine the kernels of composite systems by taking higher-order Fréchet derivatives of composite series. The generalisation of the higher-order chain rule, Faà di Bruno's formula for variational calculus, was recently determined and this note demonstrates how it can be used to determine kernels for composite Volterra series operators.

[P58] A unified approach for multi-object triangulation, tracking and camera calibration, Jeremie Houssineau, Daniel Clark, Spela Ivekovic, Chee Sing Lee, Jose Franco, IEEE Transactions on Signal Processing, Submitted October 2014 arXiv:1410.2535.

Abstract: Object triangulation, 3-D object tracking, feature correspondence, and camera calibration are key problems for estimation from camera networks. This paper addresses these problems within a unified Bayesian framework for joint multi-object tracking and sensor registration. Given that using standard filtering approaches for state estimation from cameras is problematic, an alternative parametrisation is exploited, called disparity space. The disparity space-based approach for triangulation and object tracking is shown to be more effective than non-linear versions of the Kalman filter and particle filtering for non-rectified cameras. The approach for feature correspondence is based on the Probability Hypothesis Density (PHD) filter, and hence inherits the ability to update without explicit measurement association, to initiate new targets, and to discriminate between target and clutter. The PHD filtering approach then forms the basis of a camera calibration method from static or moving objects. Results are shown on simulated and real data.

[P59] A novel approach to image calibration in super-resolution microscopy, Schlangen, I.; Houssineau, J.; Clark, D Control, Automation and Information Sciences (ICCAIS), 2014 International Conference on , vol., no., pp.111,116, 2-5 Dec. 2014.

Abstract: For many disciplines in natural sciences like biology, chemistry or medicine, the invention of optical microscopy in the late 1800's provided groundbreaking insight into biomedical mechanisms that were not observable before with the unaided eye. However, the diffraction limit of the microscope gives a natural constraint on the image resolution since objects which are smaller than half the wavelength of the illuminating light - such as proteins or ions - cannot be recognised in classical microscopy. Recently, different techniques have been developed to partly overcome this restriction using fluorescent molecules as markers. Like this, it is possible to monitor a vast diversity of intracellular processes on a molecular level which are of interest for biomedical research. Since these developments in superresolution microscopy are quite recent, suitable data analysis techniques are still to be advanced. This work aims to deploy the potential of the so-called Hypothesised filter for Independent Stochastic Populations (HISP) for multi-object estimation in a biomedical context by extending its framework to a novel joint object state and sensor drift estimator.

[P60] Multi-object filtering for space situational awareness, AAS/AIAA Space Flight Mechanics Meeting, September 2014, 15-376 Carolin Frueh; Emmanuel Delande; Daniel Clark; Jeremie Houssineau.

Abstract: This paper presents the first application to space situational awareness problems of the filter for Independent Stochastic Populations (ISP), a recent tracking algorithm derived from the novel mathematical framework for the estimation of stochastic populations, combining the advantages of traditional track-based and population-based tracking approaches. The dynamical models of Earth orbiting objects are built upon a Shepperd transition matrix and initial orbit determinations are performed based on an admissible region approach. The detection and tracking capabilities of the new filter are illustrated on a simulated five-target orbital scenario, exploiting a fixed ground-based radar.

[P61] SLAM with SC-PHD Filters: An Underwater Vehicle Application, Robotics & Automation Magazine, Chee Sing Lee; Nagappa, S.; Palomeras, N.; Clark, D.E.; Salvi, J., IEEE , vol.21, no.2, pp.38,45, June 2014.

Abstract: The random finite-set formulation for multiobject estimation provides a means of estimating the number of objects in cluttered environments with missed detections within a unified probabilistic framework. This methodology is now becoming the dominant mathematical framework within the sensor fusion community for developing multiple-target tracking algorithms. These techniques are also gaining traction in the field of feature-based simultaneous localization and mapping (SLAM) for mobile robotics. Here, we present one such instance of this approach with an underwater vehicle using a hierarchical multiobject estimation method for estimating both landmarks and vehicle position.

[P62] Sensor management with regional statistics for the PHD filter, Andrecki, M., Delande, E., Houssineau, J., and Clark, D.E., accepted in Sensor Signal Processing for Defence (SSPD), September 2015.

Abstract: This paper investigates a sensor management scheme that aims at minimising the regional variance in the number of objects present in regions of interest whilst performing multi-target filtering with the Probability Hypothesis Density (PHD) filter. The experiments are conducted in a simulated environment with groups of targets moving through a scene in order to inspect the behaviour of the manager. The results demonstrate that computing the variance in the number of objects in different regions provides a viable means of increasing situational awareness where complete coverage is not possible. A discussion follows, highlighting the limitations of the PHD filter and discussing the applicability of the proposed method to alternative available approaches in multi-object filtering.

[P63] Spawning Models for the CPHD Filter, Bryant, D.S., Delande, E., Gehly, S., Houssineau, J., Clark, D.E., and Jones, B.A., ArXiV, 30 Jun 2015.

Abstract: In the classical derivation, the CPHD (Cardinalized Probability Hypothesis Density) filter does not model the appearance of new targets through spawning. However, there are applications for which spawning models more appropriately account for new targets than birth models, with the caveat that they may create issues with computational tractability. In this paper, we present explicit formulae for the computation of the CPHD predicted intensity and cardinality distribution while accounting for spawned targets, along with three applicable spawning models. Tractability is maintained and computational complexity diminished by the use of partial Bell polynomials.

[P64] Introspective Classification for Pedestrian Detection, Blair, C.G., Thompson, J. & Robertson, N.M., September 2014, Sensor Signal Processing for Defence (SSPD 2014). Edinburgh.

Abstract: State-of-the-art pedestrian detectors are capable of finding humans in images with reasonable accuracy. However, accurate object detectors such as Integral Channel Features (ICF) do not provide good reliability; they are unable to identify detections which they are less confident (or more uncertain) about. We apply existing methods for generating probabilistic measures from classifier scores (such as Piatt exponential scaling and Isotonic Regression) and compare these to Gaussian Process classifiers (GPCs), which can provide more informative predictive variance. GPCs are less accurate than ICF classifiers, but GPCs and Adaboost with Piatt scaling both provide improved reliability over existing methods.

[P65] Event-Driven Dynamic Platform Selection for Power-Aware Real-Time Anomaly Detection, Blair, C. G., & Robertson, N. M. (2014). Video. In International Conference on Computer Vision Theory and Applications (VISAPP 2014). Lisbon, January 2014.

Abstract: In surveillance and scene awareness applications using power-constrained or battery-powered equipment, performance characteristics of processing hardware must be considered. We describe a novel framework for moving processing platform selection from a single design-time choice to a continuous run-time one, greatly increasing flexibility and responsiveness. Using Histogram of Oriented Gradients (HOG) object detectors and Mixture of Gaussians (MoG) motion detectors running on 3 platforms (FPGA, GPU, CPU), we characterise processing time, power consumption and accuracy of each task. Using a dynamic anomaly measure based on contextual object behaviour, we reallocate these tasks between processors to provide faster, more accurate detections when an increased anomaly level is seen, and reduced power consumption in routine or static scenes. We compare power- and speed- optimised processing arrangements with automatic event-driven platform selection, showing the power and accuracy tradeoffs between each. Real-time performance is evaluated on a parked vehicle detection scenario using the i-LIDS dataset. Automatic selection is 10% more accurate than power-optimised selection, at the cost of 12W higher average power consumption in a desktop system.

[P66] Identifying Anomalous Objects in SAS Imagery Using Uncertainty, Blair, C.G., Thompson, J., & Robertson, N.M., International Conference on Information Fusion, July 2015.

This paper extends the analysis in [P64] to the domain of synthetic aperture sonar. This is a more challenging modality and less data is available. However, we are able to compare performance of GPC and SVM classifiers for detecting two classes of mine-like shapes in terms of accuracy and reliability. Probabilistic SVMs outperform GPCs at this task. As we can detect multiple classes here, uncertainty information is more informative than in [P64]; we use this to detect the presence of a third class of object which the detectors were not trained on. A journal paper based on this and [P64] is in preparation.

[P67] Blair, C.G, & Robertson, N.M, 2015. Video Anomaly Detection in Real-Time on a Power-Aware Heterogeneous Platform, In IEEE Transactions on Circuits and Systems for Video Technology (accepted).

This builds on our work in [P65], extending it to include another scenario and presenting more experimental results. This knowledge of the tradeoffs inherent in choosing between FPGA and GPU has been applied to our choice of platforms to accelerate other algorithms on in WP6.1. The analysis performed here can in the future be extended to (i) other surveillance tasks and (ii) allow autonomous low-SWaP systems to automatically prioritise power consumption, detection accuracy or detection speed depending on the external conditions they are observing.

[P68] Blair, C. G., Thompson, J., & Robertson, N. M. (2015). GPU-Accelerated Gaussian Processes for Object Detection. In Sensor Signal Processing for Defence (SSPD 2015). Edinburgh, September 2015.

This extends our work on fast, reliable object detection and focuses on the use of Gaussian Process Classifiers (GPCs), which offer improved reliability and added introspective qualities compared to other detection algorithms. A major limitation of GPCs is their large computational and memory requirements, particularly in the case of dense high-dimensional data. By moving computation to GPU, optimizing the matrix processing calls which this algorithm relies on, and optimizing data arrangement and access patterns we produce a 3.7x speedup over existing BLAS-optimised code. We demonstrate this on a pedestrian detection example which was previously explored in [P64]. The techniques documented in this paper can be used when acceleration of any large computational problem relying on Radial Basis Functions (RBFs) kernels or similar is desired.

[69] Sthapit, S., Thompson, J., Hopgood, J. & Robertson, N (2015). Distributed Implementation for Pedestrian Re-identification in Sensor Signal Processing for Defence (SSPD 2015). Edinburgh, September 2015.

Pedestrian re-identification in distributed systems allows a person to be tracked and re-identified as they move from one non-overlapping camera viewpoint to another. An important stage in this task is the algorithm used for feature extraction and comparison. For battery-powered processing nodes used at each camera (such as mobile phones), the computational requirements and the volume of data transferred between nodes must be minimised to allow lower power consumption and hence longer battery life. Here we compare three key algorithms and evaluate their discriminative ability. We show that the Keep It Simple and Straightforward Metric (KISSME) provides the best balance between discriminative performance, computational complexity and feature length.

References

[R1] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, “Beyond Nyquist: Efficient Sampling of Sparse Bandlimited Signals”, IEEE Transactions on Information Theory, vol. 56, no. 1, pp. 520–544, 2010.

[R2] M. Mishali and Y. C. Eldar, “From Theory to Practice: Sub-Nyquist Sampling of Sparse Wideband Analog Signals”, IEEE Journal of Selected Topics in Signal Processing, vol. 4, no. 2, pp. 375–391, 2010.

[R3] P. Feng and Y. Bresler, “Spectrum-blind minimum-rate sampling and reconstruction of multiband signals”, IEEE International Conference on Acoustics, Speech, and Signal Processing, 1996. ICASSP-96. Conference Proceedings, 1996, vol. 3, pp. 1688–1691 vol. 3.

[R4] C.H. Gierull, “Digital Channel Balancing of Along-Track Interferometric SAR”, Technical Memorandum, DRDC Ottawa, 2003.

[R5] S. R. J. Axelsson, “Position correction of moving targets in SAR Imagery”, Proc. SPIE, vol. 5236, pp. 80-92, Jan. 2004.

[R6] M.P. Dana, “Registration: A prerequisite for multiple sensor tracking,” in Multitarget-multisensor tracking: Advanced Applications, Chp.5, pp. 155—185 , 01/1990.

[R7] N. Kantas, A. Doucet, S. Singh, and J. Maciejowski, “An overview of Sequential Monte Carlo methods for parameter estimation on general state space models,” in Proceedings of the 15th IFAC Symposium on System Identification, 2009, pp. 774—785.

[R8] E.B. Sudderth, A.T. Ihler, M. Isard, W.T. Freeman, and A.S. Willsky, “Nonparametric Belief Propagation,” Communications of the ACM, vol. 53, no. 10, pp. 95—103, Oct. 2010.

[R9] B. Ristic, D. Clark, B.-N. Vo, B.-T. Vo, “Adaptive target birth intensity for PHD and CPHD filters,” IEEE Transactions on Aerospace and Electronic Systems, vol. 48, no. 2, pp. 1656—1668, 2012.

[R10] R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter, Probabilistic Networks and Expert Systems. Springer, 1999.

[R11] N. Patwari, J.N. Ash, S. Kyperountas, A.O. Hero, R.L. Moses, N.S. Correal "Locating the nodes: cooperative localization in wireless sensor networks," IEEE Signal Processing Magazine, vol.22, no.4, pp.54,69, July 2005.

[R12] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, “Measuring and testing dependence by correlation of distances,” *The Annals of Statistics*, vol. 35, pp. 2769–2794, 2007.

[R13] C. Piciarelli, and G.L Forest. “On-line trajectory clustering for anomalous events detection”. Pattern Recognition Letters, 27(15), pp 1835-1842, 2006

[R14] Statistical Multisource-Multitarget Information Fusion, Mahler, R.P.S., Artech House, 2007.

[R15] Representation and estimation of stochastic populations, Houssineau, J., PhD Thesis, 2015.

[R16] Multiple Hypothesis Tracking for Multiple Target Tracking, Blackman, S.S., IEEE Aerospace and Electronic Systems Magazine, Jan 2004.

[R17] Calibration of Multi-Target Tracking Algorithms Using Non-Cooperative Targets, Ristic, B.; Clark, D.E.; Gordon, N., Selected Topics in Signal Processing, IEEE Journal of , vol.7, no.3, pp.390,398, June 2013.

[R18] SLAM With Dynamic Targets via Single-Cluster PHD Filtering, Chee Sing Lee; Clark, D.E.; Salvi, J., Selected Topics in Signal Processing, IEEE Journal of, vol.7, no.3, pp.543, 552, June 2013.

Abstract: This paper presents the first algorithm for simultaneous localization and mapping (SLAM) that can estimate the locations of both dynamic and static features in addition to the vehicle trajectory. We model the feature-based SLAM problem as a single-cluster process, where the vehicle motion defines the parent, and the map features define the daughter. Based on this assumption, we obtain tractable formulae that define a Bayesian filter recursion. The novelty in this filter is that it provides a robust multi-object likelihood which is easily understood in the context of our starting assumptions. We present a particle/Gaussian mixture implementation of the filter, taking into consideration the challenges that SLAM presents over target tracking with stationary sensors, such as changing fields of view and a mixture of static and dynamic map features. Monte Carlo simulation results are given which demonstrate the filter's effectiveness with high measurement clutter and non-linear vehicle motion.

Contact Information

Academic Staff

Prof Mike Davies, Director, Work Package 1 Leader mike.davies@ed.ac.uk

Prof Yvan Petillot, Deputy Director, Work Package 3 Leader y.r.petillot@hw.ac.uk

Prof Bernard Mulgrew, Work Package 2 Leader Bernie.Mulgrew@ed.ac.uk

Dr Neil Robertson, Work Package 4 Leader N.M.Robertson@hw.ac.uk

Dr Daniel Clark, Work Package 5 Leader d.e.clark@hw.ac.uk

Prof John Thompson, Work Package 6 Leader john.thompson@ed.ac.uk

Prof Andrew Wallace, Work Package 3 & 6 a.m.wallace@hw.ac.uk

Dr James Hopgood, Work Package 3 & 4 James.hopgood@ed.ac.uk

Dr Mathini Sellathurai, Work Package 2 M.Sellathurai@hw.ac.uk

Project Management

Janet Forbes, Project Manager janet.forbes@ed.ac.uk