April 2016 to March 2017

UDRC – Year 4 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, Cardiff University and Newcastle University (LSSCN Consortium).

This work is funded by the MOD and EPSRC.

Vision

The future battlespace will be a complex environment characterised by known and unknown threats, modern and legacy sensor systems, a congested RF spectrum, and mobile and static forces. Information is key in warfare but future conflicts are likely be characterised by an increased level of complexity in intelligence gathering and analysis. Unless such complexity can be overcome, the effectiveness of critical decision making and operational actions will be reduced. Furthermore, many of these issues are mirrored in civilian contexts and therefore the contributions from the programme will have wider impact.

Historical, current and future sensor systems will provide ever more data for subsequent analysis, hence advances in technology will be essential to ensure that they can be optimally exploited. The outputs of sensors of different modalities, capabilities and locations within the battlespace will need to be combined in multiple ways so that such optimal exploitation can be ensured in a wide variety of operations at all levels of conflict. However, at the same time, the electronic environments in which such conflicts will take place are likely to pose greater problems as the availability of bandwidth becomes ever more restricted.

Given the significant operational and technical challenges outlined above, our desired outcome for the research programme is to identify techniques and technology that will increase the situational awareness of our fighting forces to a level that will represent a significant increase in the probability of mission success. This will be achieved, inter-alia, by the efficient, effective and timely processing and communication of the wide range of available sensor data

Specifically, we will provide transformational new signal processing solutions which exploit multi-sensor and multimodal data, whilst retaining bandwidth and computational efficiency, to maximize the UK's defence capabilities and its broader academic and industrial skill-base in signal and data processing. In particular, we believe that networked-enabled distributed sensing should provide new capabilities such as combating stealth. However, this potentially increases the complexity of the processing task. This could be mitigated by new signal-separation/beamforming algorithms utilising sparsity concepts. Control and management of distributed sensors could be costly in terms of network traffic so systems that are able to interact without central control would be preferable. Finally, in order to protect this new networked-enhanced sensing paradigm, aspects of cyber-security will play an important role. A measure of our success will be the extent to which our signal processing solutions have enhanced the defence technology base.

In addition to this, we aim to build a healthy community of practice in defence signal processing spanning academia, industry and government. Through this, we envisage the emergence of the next generation of signal processing engineers to strengthen the UK's leading position in this area.

Staff and Students

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

The following research groups are involved in this work:

- Institute for Digital Communications (IDCOM), University of Edinburgh.
- Oceans Systems Laboratory (OSL), Heriot-Watt University.
- Signal and Image Processing Laboratory (SIP-Lab), Heriot-Watt University.
- School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast.

Academic Staff

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

Research Associates

Dr Murat Uney
Dr Yan Pailhas
Dr Rolf Baxter
Dr Emmanuel Delande
Dr Cris Rusu
Dr Salvatore Caporale
Dr Steven Herbert
Dr Christian Nunez Alvarez
Dr Paulo Garcia

Research Students

Di Wu
Puneet Chhabra
Jose Franco
Saurav Sthapit
Kimin Kim
Alessandro Borgia
Alexey Narykov
David Cormack

Project Management

Janet Forbes
Sebastien Georges
Audrey Tainsh
Maddy MacBeath

Research

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.

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, Localisation, 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 Estimation Framework for Multi-target Detection/Tracking and 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

Highlights

WP1 Sparse Representations and Compressed Sensing

a) Extending the Volumetric SAR data calibration algorithm to the short aperture processing mode. This mode can:

1) better calibrate the collected raw data in a multipath trial and;
2) fully explore the capability of the proposed calibration technique.

We managed to handle such a short aperture SAR processing, which is computationally infeasible using ad hoc exhaustive searches [P1].

b) The nonlinearities in Raman spectral mixture decomposition have been explored, where these nonlinearities come in the form of global spectral shifts, due to imperfect spectroscopy, and/or local spectral shifts, due to complex interactions between chemical bonds. The compensation of such artefacts was investigated and novel techniques were proposed to model such behaviours, to refine spectral decompositions [P2].

c) Advanced joint fluorescent background removal and Raman spectral decompositions. This idea is feasible due to morphological differences between the signals, i.e. background and Raman signals. A joint spectral decomposition using an augmented library was proposed and some preliminary results were reported [P2].

d) Refined SAR GMTI algorithm to guarantee the convergence of iterative algorithm. While the original algorithm practically performs well, it was not clear if the algorithm converges in different settings of sparse SAR GMTIs. Some modifications were proposed to ensure the exponential convergence of the algorithm, which will be reported in the revised version of [P3].

WP2 Distributed Multi-Sensor Processing

a) Developed a coherent long time integration algorithm for detection of manoeuvring dim targets with mono-static [P11], bi-static [P12], and, multi-static radar configurations [P13] that occur in a distributed radar system.

b) Developed an adaptive waveform design approach for active multiple input multiple output (MIMO) sensor systems that minimises estimation errors in a multi-target environment [P14]. This algorithm can be used for both distributed and co-located configurations.

c) Investigated asymptotic properties of approximate likelihoods which were previously introduced for scalable parameter estimation in multi-sensor multi-target tracking models, and, developed an empirical Bayesian approach for their use in a parameterised general multi-target model [P15].

WP3 Unified Detection, Localisation and Classification

a) Participated in the the ONMEX16 and MANEX16 trials organised by CMRE in September/October 2016 allowing us to collect 2.8TB of wideband sonar data on diverse seabed types and for various mine-like objects and be used to study acoustic coherence [P24].

b) Theoretical derivation of the CSAS (Circular Synthetic Aperture Sonar) PSF (Point Spread Function) and CSAS image deconvolution. [P17, P21, P23]. Circular trajectories offer interesting problematics in terms of imagery for SA. We analytically derived the Point Spread Function (PSF) for such systems and proposed a new processing algorithm to improve the resolution.

c) Design, implementation and simulation of a novel SAS micronavigation algorithm allowing for improved focusing and cheaper SAS operations. An experiment has been performed on a SAS system proving its effectiveness, accuracy and robustness on real-data [P25 ,P35].

d) A fast kernel discriminatory orthogonal Dictionary Learning (DL) method has been proposed to classify large-scale (>60,000 samples) and high-dimensional (>90000) datasets. State of the art classification accuracy of 97% is reported when compared with benchmark datasets and Deep Learning algorithms. Under similar experimental settings ~55 times run-time speed up has been achieved. This has been submitted to EUSIPCO 2017 [P28].

e) For Lidar we have proposed a new unsupervised full-waveform and peak classification algorithm that simultaneously extracts peaks (which results in a 3D point cloud) and classifies them. No assumptions on the target classes is made [P33].

WP4 Context Driven Behaviour Monitoring and Anomaly Detection

a) Initial work on unifying pattern-of-life learning, anomaly detection and tracking, focusing on familiarisation with the Hypothesised filter for Independent Stochastic Populations (HISP) [R10] and with fully probabilistic methods for pattern-of-life modelling and the combination of a modified HISP filter with online Kernel Density Estimation [R11] (oKDE) has shown that a compact probability density function can be constructed to model spatial context in an online setting. oKDE has shown that it can model the true spatial distribution of target pattern-of-life with only 12% error while compressing the data that must be retained by 97%.

b) Implemented a baseline framework for our work on people tracking by re-identification which performs convenient pre-processing of the two largest benchmark datasets and is able to extract discriminative deep features by a residual learning-based Convolutional Neural Network (CNN).

c) The Temporal Anomaly Detection enabling contract addressed the problem of anomalous device identification in noisy broadband environments. This work showed that the Symbol Aggregate approximation (SAX) algorithm [R12] performs poorly in such environments, with weak anomalous devices remaining undetectable with four significant outputs:

1. A novel algorithm for removing noise from spectrogram data;
2. Frequency-SAX: An extension of the SAX algorithm to incorporate frequency-based features that permits the detection of transient unknown devices;
3. SEGM: A new image segmentation based approach which performs automatic clustering of spectrogram images enabling anomalous device detection and re-identification;
4. Implementation code for each algorithm has been supplied to Dstl

WP5 Estimation Framework for Multi-target Detection/Tracking and Sensor Management

a) Submission of a comprehensive description of the general structure of DISP filter, including an adaptation of the information metric for sensor management [R18] allowing for the exploration of the state space for yet-to-be-detected targets [P41].

b) Implementation of a filtering framework in native C++ code for the sake of efficiency, with I/O formatted in XML files to allow for the comparison of different filtering solutions (either developed in our research group or externally) on the same scenario. Integration of the DISP filter to this framework.

c) Implementation of the HISP filter, an approximation of the DISP filter for large scale scenarios, and which has been successfully adapted to the context of Space Situational Awareness (SSA) [P48], where the number of satellites and orbital debris is one of the main challenges to be addressed for the construction of a catalogue of orbiting objects.

WP6 Efficient Computation of Complex Signals Processing Algorithms

a) Developed computationally efficient algorithms which focuses on a low computational complexity, allowing them to run on a hardware platform with limited computational and power limited capabilities such as mobile devices or other wearables.

b) Formulated the sensor management problem as a binary optimization problem which was solved by convex relaxation methods and uses infinity regularized reweighted L1 penalty functions in an iterative fashion to create an activation schedule for the whole sensor network over all time instances. This method is computationally efficient to solve (polynomial complexity)[P53].

c) Computational offloading in sensor networks [P52]. With the goal of balancing energy consumption in a sensor network, we propose to offload some processing that takes place at a sensor node to neighbouring sensors. In this fashion we improve energy efficiency and increase the robustness of the sensor network by making sure the processing that takes place at the sensors do not deplete their energy, putting them offline. This line follows research efforts to include in computational models and also edge computing.

WP1 Sparse Representations and Compressed Sensing

Research Leader: Mike Davies

Academics: Mehrdad Yaghoobi, Bernard Mulgrew, Mathini Sellathurai, John Thompson, Yvan Petillot

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 the 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 worthwhile for further investigations. 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 the ultra-wide band Radio Frequency (RF) signal conversion.

Compressed sensing is incorporated to present a more efficient method 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. In WP1.3, we extend compressive sampling to the settings, in which the task is not the reconstruction, but some sort of inference. In this part, we explore the capability of undersampled signal processing to extract figures, like the contribution of each factor, classify the signal or quantify the inputs. The focus of this sub-workpackage would be on moving target localisation in multichannel SAR and Raman unknown mixtures fingerprinting and quantification.

Outcomes

The multipath volumetric SAR imaging was investigated by presenting a systematic way of calibration of the pulse information, collected in different passes. Such a calibration is essential for high resolution volumetric imaging. Such a data has usually been calibrated using reference targets and some domain knowledge about the locations and antenna profiles. Such information may not be available in practice, which justifies a systematic approach to auto-calibrate the data. The calibration problem was formulated as a phase retrieval problem for the first time and solved using an adapted phase recovery technique[P4,P5,P7]. The new method can use a scene with one or some bright targets and automatically find the range estimation error to have a sharp volumetric image of the scene. The algorithm can incorporate short to long apertures to correct the range estimation errors, caused by inaccuracy in antenna locations.

The SAR GMTI was reformulated as the separation of sparse moving targets from static background. Such formulation allows us to separate these components using a multi-channel SAR system. We demonstrated that two channels are not enough to separate moving targets from static background. We then presented a practical method to solve the corresponding optimisation program [P4,P6,P10]. Recently, we demonstrated that a variant of this algorithm asymptotically converges. It was tested with the GOTCHA dataset, having three receive channels.

A fast spectral decomposition technique was developed for Raman spectral mixtures, for running on an embedded platform. The algorithm is based on the fast non-negative sparse representation of the spectra using a reference library and by indicating a possible out of library element [P9, P8]. The computational cost and memory usage of the algorithm are low and it is suitable for implementation on limited computational power platforms, like handheld devices. The algorithm attracted the interests of Dstl and variations of the algorithm have been developed for the spectral decomposition with various levels of mixtures, nonlinearities and background contaminations [P2]. The low-level language implementations of the algorithm is now available in C and Java for importing in to the handheld devices.

Progress

The research on volumetric SAR imaging had significant progress in the data calibration part. Raw data from multiple trials, recorded for a full volumetric imaging, normally has some range estimation errors, due to the inaccuracy of navigation units and other source of errors. To accumulate pulse information, we need a pre-processing step to fix the errors. The error correction task was formulated as a phase retrieval problem, where the phase in Fourier domain translates to the spatial shifts. Finding the correct phase was done using a modified canonical method, called the Gerchberg-Saxton technique, which is tailored for the settings of this problem. Tomographic view of a civilian car before and after range focusing has been shown in figure 1 [P1].

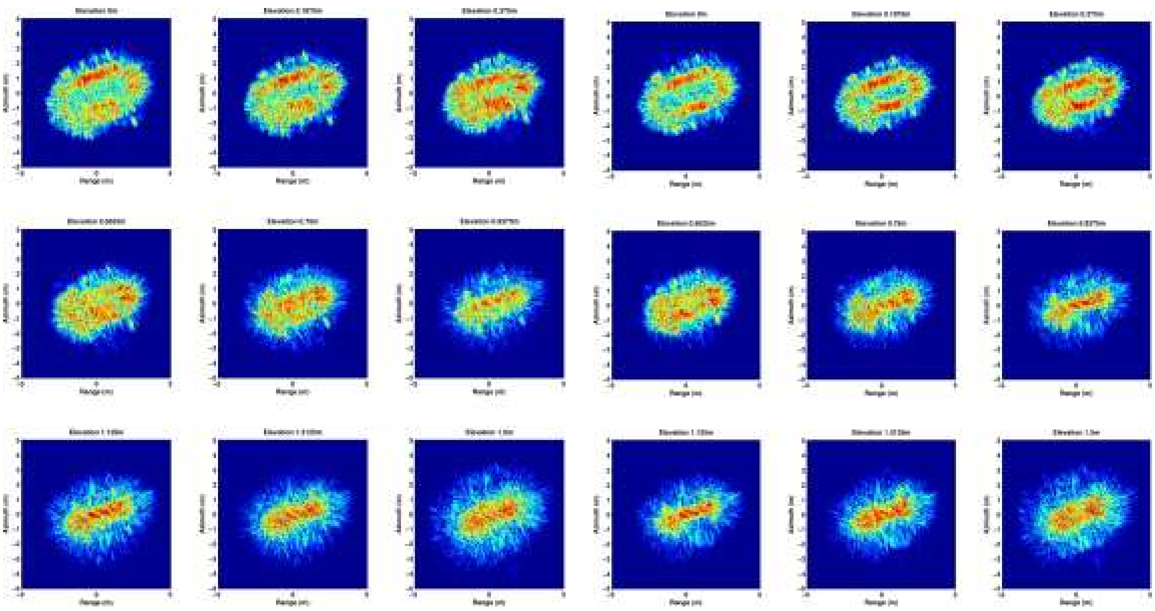

Figure 1: Sliced tomographic view of the reconstructed volumetric image of mcar3 in different heights, using the original (left) and the corrected (right) phase histories.

The moving target tracking in SAR, which was formulated as a sparse approximation problem, refined using digital elevation map of the scene. We use the backprojection with the incorporation of DEM to accurately reconstruct the image and estimate the moving targets. When the scene is not flat, we estimate the normal velocity to refine overall estimations. The proposed sparsity SAR GMTI algorithm with DEM and then dimentional velocity estimation has been applied to GOTCHA SAR GMTI data set and the location of target and the velocities are plotted in figure 2. The new algorithm managed to estimate the location and velocities more accurately, with an expense of more computation [P3].

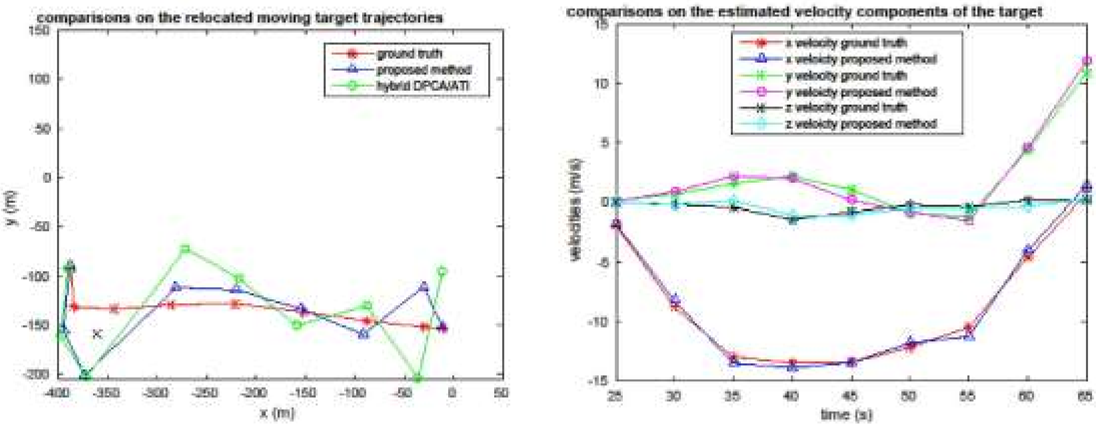

Figure 2: The comparisons of relocated target path and target velocity components.

The Raman spectral decomposition for chemical mixture fingerprinting and quantification was exploited by considering nonlinearity artefacts, in the form of global and local spectral shifts. The global shifts are mainly due to imperfect instrumental measurements. The local shift of Raman peaks are due to complex chemical bonds interactions in some mixtures. These two can cause misdetection of some chemicals in the mixtures and/or incorrect quantifications. Some methods were tested to compensate such artefacts [P2], which need more investigations in the future. Raman spectral decomposition needs a fluorescent background separation process, which has usually been done in the preprocessing step. We started investigations of the joint background separation/Raman decomposition. The potential benefit of such an approach is to stop propagation of the background removal errors to the decomposition step. The ideal scenario is to remove the background while decomposing the spectra to the elementary components [P2].

Future Direction

Volumetric LF-SAR: the volumetric compressive SAR imaging will be applied to the wideband low-frequency SAR. The penetration effects of LF radio frequency signals need to be considered to yield a more accurate sliced tomographic view of the objects. LF-SAR has to be notched to stop interfering with the occupied frequency bands. Such an interruption in the frequency bands generates large side-lobes, which distorts the image. The interrupted SAR imaging can be formulated as a missing data problem and be (approximately) solved using a compressed sensing framework, which we will explore in the next year of this project.

Non-linear Raman Spectral Decomposition: the linear model is sometimes not accurate for the spectral decomposition, particularly with the small contributions of chemicals. We explore a new framework for spectral shift modelling, using a first order approximation of the nonlinearity [P2]. In this framework, an augmented library, including extra elements related to the first order approximations, is fed to a sparse approximation method. Theoretical and practical aspects of such an approximation have to be investigated.

Research Leader: Bernard Mulgrew

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

Research Associate: Murat Uney, Steven Herbert

PhD Student: Kimin Kim

WP2 Distributed Multi-Sensor Processing

The objective of this work package is to address 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.

The second stage of this research, WP2.2 Distributed/Decentralised detection, was started in 01/2015 following the first sub-workpackage WP2.1 Fusion and registration (06/2013-12/2014). In this research, we have been addressing challenges in detection of targets with networked sensors by exploiting the diversity and/or extended coverage provided by having more than one sensor. Detection of manoeuvring and/or dim targets is particularly challenging. We explore 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

This research thrust has resulted in an algorithmic framework for coherent detection of manoeuvring dim targets that captures mono-static, bi-static and multi-static configurations with phased array radar receivers [P11, P12, P13]. These configurations arise in multiple radars with non-synchronous orthogonal waveform transmitters. We have introduced a statistical inference algorithm which simultaneously synchronises separated tx/rx channels, estimates trajectories and integrates signal returns along these trajectories for use in a Neyman-Pearson test.

We also addressed waveform design for active multiple-input multiple output (MIMO) sensory systems by establishing a general objective function to optimally design probing signals that achieve best reduction in uncertainty with a given transmission energy budget. Our approach can be used for both co-located and separated transmitters and is demonstrated in simulations for a simple MIMO radar system [P14].

The third line of investigation is an extension of our previous work on approximate likelihoods for scalable parameter estimation in multi-sensor state space models [R1, R2]. We aimed to employ a relatively more accurate approximation given in [P16] in a multi-target environment, and, developed an empirical Bayesian approach within a parametric general multi-target tracking model for this purpose. The result is a distributed sensor self-calibration algorithm for fusion networks [P15].

Progress

In this 12 month period, the main line of our research activities has been aligned with WP2.2 Decentralised/Distributed Detection and the UDRC Future Strategy detailed in the UDRC mid-term review report. In particular, we have been developing models and algorithms for detection of dim and manoeuvring targets with multiple transmitters and array receivers (for an example cooperative system, see Figure 3). We have further extended these efforts towards addressing the problem of waveform design in active sensing from a general perspective that can address a variety of multi-input multi-output (MIMO) configurations.

In addition, we have worked on extending our outcomes in WP2.1 Distributed Fusion and Registration which introduced an approximation framework for scalable distributed estimation in fusion networks. We improved both our theoretical understanding and the applicability of our approximations, and, developed methods for their use with different multi sensor multi target tracking models.

For detection in active sensor networks, we have developed models and algorithms for maintaining coherence across parts of a network using the target trajectory as a reference [P11, P12, P13]. First, we tested our modelling approach in a mono-static configuration with a co-located transmitter/receiver, and, demonstrated that it is possible to perform long time integration while simultaneously estimating the target trajectory to take target manoeuvres into account in a dynamic programming fashion [P11] (also, see Figure 4). In [P12], we extended this approach for separated transmitter/receiver pairs, i.e., bi-static channels, with an unknown time reference shift. In particular, we recover the synchronisation term by diverting simultaneous beams towards the tested point of detection and the remote transmitter, thereby relaxing the commonly used assumption that the remote transmitters and the local receiver are synchronised (see, e.g.,[R3, R4]). In [P13], we introduce a novel solution for the multi-static case in which both mono-static and bi-static channels exist. Overall, this work addresses processing at one of the receivers in Figure 1 using probing signal emission of any of the transmitters.

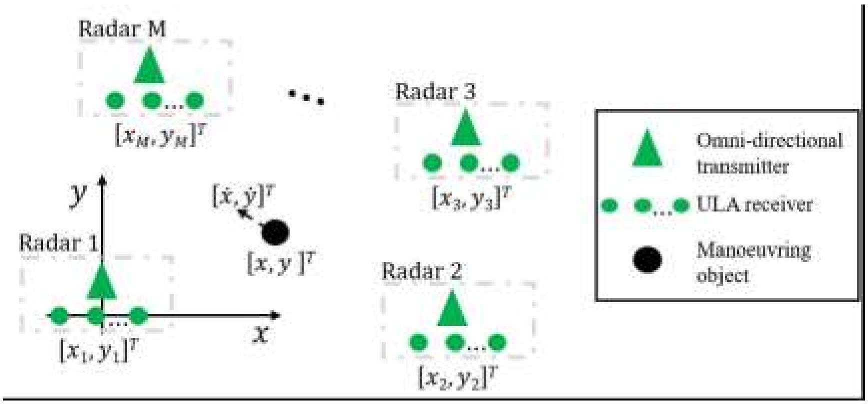

Figure 3. A radar network as a cooperative active sensing system: Each node has an array receiver and is capable of spatial filtering (through beamforming). The transmitters use orthogonal waveforms and their transmission characteristics are fully known to other nodes in the network up to an unknown time reference shift (i.e., non co-located tx/rx pairs are not synchronised).

In a more general perspective, we considered active MIMO sensing that can have both co-located and separated tx/rx configurations and the problem of designing probing waveforms for the transmitters. In [P14], we have achieved the aim of expressing a general cost function to optimise for minimum mean squared error adaptive waveform design, for estimation of parameters associated with a fixed, known number of targets. Previously only approximate cost function expressions existed. Furthermore, we have proven the principle of our method using multiple targets, whereas only a single target has been considered in the existing literature [15].

An extension of our work in WP 2.2 (e.g., [R1, R2]) appeared in Fusion 2016 [P16] which addresses multi-sensor calibration in fusion networks with sensors that have partially overlapping field-of-views. We further developed an empirical Bayesian approach to use the approximations proposed in [R1, R2] in parameterised multi-target tracking models (see, e.g., [R6]). This approach and the resulting self-calibration algorithm is reported in detail in [P15].

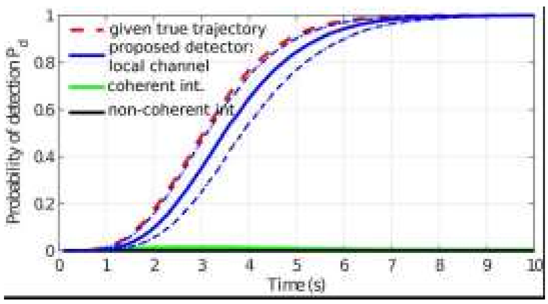

(a)

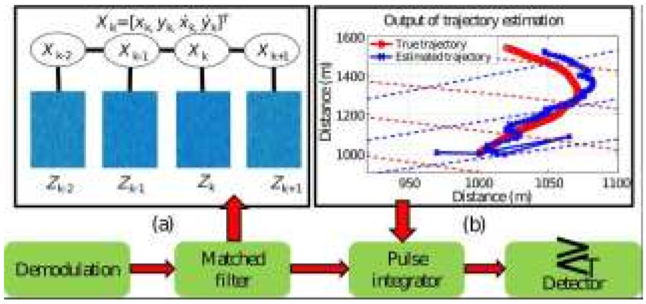

(b)

Figure 4. (a) The block diagram for simultaneous trajectory estimation and long time integration. (b) Example performance given by the integration time versus probability of detection of this approach (blue line) in comparison to integration with the full knowledge of the true trajectory (dashed red line) and conventional coherent and non coherent integration (green and black lines).

Future Direction

The long time coherent integration work will be extended to accommodate the capability of simultaneously estimating signal features such as micro-doppler signatures. This will underpin the capability of extracting signatures for manoeuvring targets in adverse background conditions and lead to their identification.

We also aim to address cooperation at the receivers in a radar network through design of efficient communication strategies for improving local detection accuracy. We will be extending the current work on distributed detection/estimation (e.g., [R7, R8]) to models that capture active sensing in a multi-target environment (e.g., [R9]).

The adaptive waveform design research will continue to investigate whether our cost function can be better optimised compared to our current solution as it is a non-convex optimisation problem. One particular goal will be the generalisation of this cost function expression for the case where we do not assume that the number of targets is fixed and known.

Research Leader: Yvan Petillot

Academics: Daniel Clark, James Hopgood, Andrew Wallace

Research Associates: Yan Pailhas, Salvatore Caporale

PhD Student: Puneet Chhabra

WP3 Unified Detection, Localisation and Classification in Complex Environments

The aim of this work package 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. Therefore 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

Our successful participation to the CMRE NATO trials in September/October 2016 allowed us to collect around 2.8TB of wideband sonar data on diverse seabed types and for various mine-like objects. The dataset has been collected specifically to study acoustic coherence. An unclassified report describing the equipment, the sonar specifications, the methodology and the data is now available.

With regard to SAS micronavigation, we proposed a novel algorithm not requiring any constraint nor any prior knowledge on the motion trajectory. The technique has initially been shown to be capable of estimating subwavelength motions on synthetic data and then successfully tested on real data acquired in a tank by means of a real SAS system. Further achievements have also been obtained on theoretical aspects related to the computation of the backprojection required for the SAS image formation.

Progress

From 14th September 2016 till 6th October, Yan Pailhas and Nicolas Valeyrie participated to the ONMEX’16 and the MANEX’16 trials (fig. 5-8). The two trials were organised by CMRE, and took place respectively in the Bay of Hyeres, close to Toulon in France and in Framura in Italy. The wideband data was collected with the Hydrason BioSonar UWBS (Ultra WideBand Sonar) array. In total we performed more than 20 missions. REMUS traveled over 175km inspecting around 13km^2 of seafloor. The general scope of these trials was to collect a substantial data set using the WBMBS (Wideband Multi-Beam Sonar) system and three sidescan sonars to study the problem of coherence. The WBMBS has a very broad band of frequencies (20-180kHz exploitable). It also has a very wide beam pattern (40deg @ 60kHz). In a similar way to SAS systems, the WBMBS ”sees” every particular

point in the scene numerous times. It is therefore particularly well adapted to measure spatial coherence. The multi-element aspect of the WBMBS allows us, via adaptive processing, to maximise the SRR (Signal over Reverberation Ratio), and thereby to track a cleaner measurement of the coherence of a particular point in the scene. We aimed to carry out repeated measurements at different grazing angles and different aspects in a number of different types of environment to assess the limit of coherence loss and its dependency on look-angle, frequency, seabed type, etc. A special emphasis has been on any man-made targets present in the environment and polygonal/circular target re-acquisition have been performed to maximise information gain. The output of the trials will deliver a necessary data set to answer fundamental questions about coherence as well as material to develop recognition algorithms based on coherence processing.
In the framework of SAS micronavigation, state-of-the-art approaches rely on accurate inertial trajectory control. Conversely, the proposed technique relies only on the observation operator, thus can be employed in much less restrictive conditions [P25, P35]. An experiment has been performed in a tank by using a real SAS system provided by Hydrason whose results are shown in Fig. 9. Despite many non-idealities, the motion trajectory has been estimated with a subwavelength accuracy and the method has been revealed to be robust.

With regard to theoretical signal processing for underwater environments, a novel approach for compensating for the acoustical speed variation at low frequency has been proposed in cooperation with ENSTA, Bretagne [P26]. By taking advantage of the group delay shift covariance property of frequency warping transformations, the energy profile relative to each propagating mode can be isolated, thus increasing the knowledge on the propagation mean and potentially resolving more information about the acoustic signal source.

As far as numerical techniques for synthetic aperture processing are concerned, a theoretical study on the computation of a special class of non-uniform Fourier transform over the time domain (as it is needed for SAS imaging) has been pursued. The outcome comprehends an accurate algebraic modeling of both the direct and the inverse transform and the corresponding decompositions and factorizations allowing for fast computation [P27, P36].

The associated PhD: Anomaly Detection and Object Classification using Multi-spectral LiDAR [P29, P31, P32, P33] and Wider-band Sonar (Puneet Chhabra) [P31] will be submitted shortly. This work focusses on signal processing algorithms applied to full-waveform LiDAR and wide-band sonar and outputs are as follows:

a. Improved Image Discrimination using Fast Kernel DL [P28]

Most real-world signals or images have an intrinsic non-linear similarity measure and can be harder to discriminate. Kernel dictionary learning with applications to signal classification offers a solution to such a problem. However, decomposing a kernel matrix for large datasets is a computationally intensive task. Existing papers on dictionary learning using optimal kernel approximation method improve computation run-time but learn an over-complete dictionary. We show that the learning and classification run-time can be significantly decreased if we learn a discriminative orthogonal dictionary instead. The proposed algorithm, Kernelised simultaneous approximation and discrimination (K-SAD) [P28], learns a single highly discriminative and incoherent non-linear dictionary on small to medium-scale real-world datasets (RGB-D and face datasets, see [P28] for more details).

Extensive experiment [P28] results in ~97% classification accuracy on publicly available datasets. We do not extract any features on the image data and use the raw data as an input to our algorithms. However, when we apply this technique to LiDAR and Sonar data for signal discrimination minor modifications are necessary. For example, raw full-waveform lidar histograms need to be processed to extract individual peaks.

b. Aerial and Terrestrial LiDAR

Alongside underwater multi-spectral single photon counting lidar experiments as part of the Puneet Chhabra's PhD work, we have carried out terrestrial and aerial scanning using COTS LiDAR sensor. This work is on-going in collaboration with Carbomap Ltd, UK and Riegl GmbH, Austria. As part of the surveying a back-pack system (Figure 10) was developed that could carry a monochromatic LiDAR for data collection. University of Edinburgh's Kings Building campus was scanned on several occasions under different conditions. We have developed a peak detection and classification technique that allows to generate a 3D point and classify them simultaneously. Figure 11 illustrates a snapshot of the unsupervised classification results on the Austria LiDAR datasets. The data was collected at an altitude of > 300 meters on two separate occasions using two different sensors operating at different wavelengths. We process the full-waveforms and build the 3D point cloud and classify them simultaneously. Finally, Figure 12 illustrates terrestrial scans of Kings Building campus using the backpack system shown in Figure 10. The colors correspond to intensity which his shaded for vegetation, grass, buildings and man-made surfaces.

Figure 5: (left) CMRE's research vessel: the ALLIANCE, (right) REMUS equipped with the UWBS system on the dock of the small boat ready to be deployed.

Figure 6: Examples of mine-like objects been deployed in the area of interest.

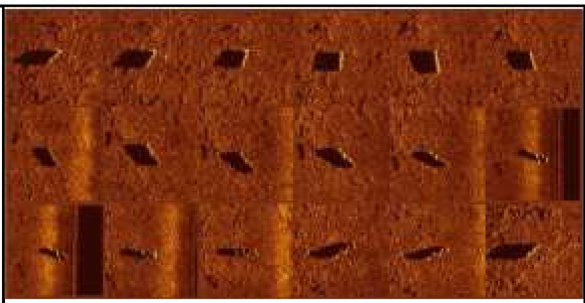

Figure 7: Discrete re-acquisition pattern of one of the mine-like object.

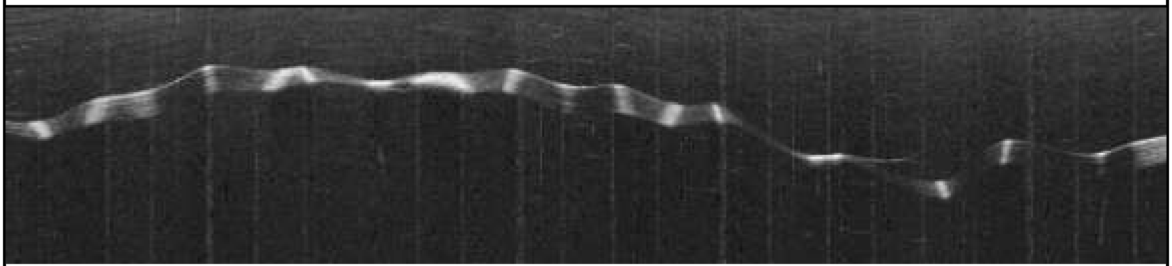

Figure 8: Snapshot of the continuous acquisition of one of the mine-like object.

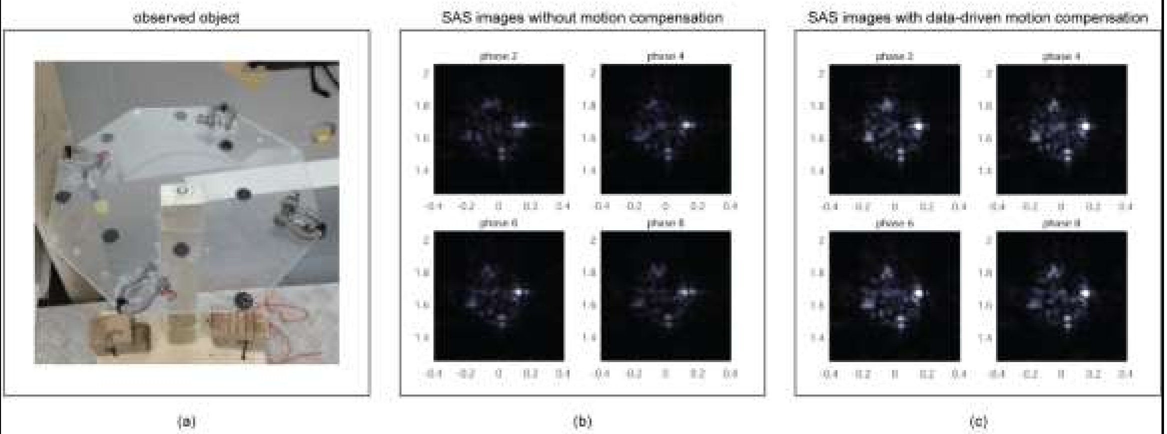

Figure 9: Test object consisting of four strong reflectors (a), recovered reflectivities for different placements of the SAS pings with subwavelength motion errors (b) recovered reflectivities with compensated motion errors by means of the proposed novel technique. The fact that the images in (b) are visibly different implies the presence of motion errors, whereas the likelihood of the images in (c) implies that the motion has been correctly estimated and compensated.

Figure 10: A backpack system developed for data collection using Riegl's VUX-1LR Lidar sensor.

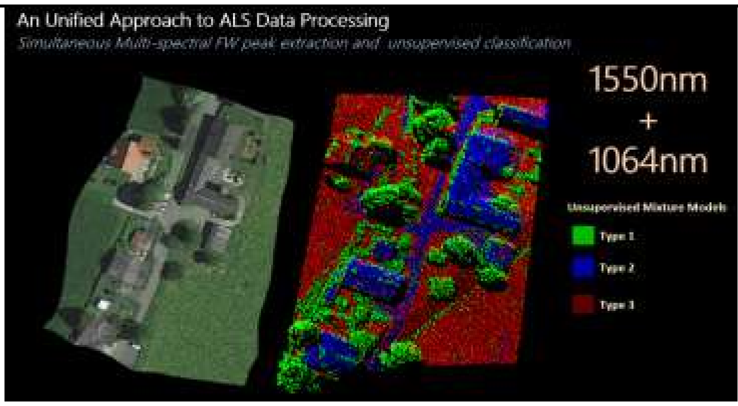

Figure 11: Simultaneous peak-extraction and classification of aerial full-waveform LiDAR.

Figure 12: Kings Building scans using VUX-1LR and custom backpack system developed by us.

Future Direction

Over the next year, we will focus our effort on the study of the acoustical coherence. Thanks to the data collected during the ONMEX'16 and MANEX'16 trials, a large dataset is now available for theory validation and early algorithms testing. Part of our effort will also be dedicated to investigating further the problems inherent to SAS (Synthetic Aperture Sonar) imaging and SAS processing.

In consideration of the positive outcome of the proposed micronavigation technique on real SAS systems, follow-on experiments will be planned to assess potential employments in conditions where state-of-the-art systems are not able to operate. Moreover, the theoretical principle supporting this methodology could be exploited in different scenarios such as MIMO and interferometry.

Interactions

We have recently been awarded an EPSRC research project on acoustic sensor networks (USMART) in collaboration with Newcastle and York Universities for a total value of £1.2M (HWU share 448K) aiming towards developing a generic and affordable acoustic sensor network for sensing with applications to the oil & gas industry, marine science and defense. Details are available here: http://gow.epsrc.ac.uk/NGBOViewGrant.aspx?GrantRef=EP/P017975/1

WP4 Context Driven Behaviour Monitoring and Anomaly Detection

Research Leader: Neil Robertson

Academic: 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. 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 using spatial context acquired through pattern-of-life learning from wide-area surveillance.

As a subcomponent of this, we are also addressing the problem of people tracking in networks non-calibrated, non-overlapping cameras. Our aim is to define a strategy for tackling long occlusions by proposing a unified framework for tracking with deep-learning-based target re-identification. An iterative adaptive interaction is created between these two tasks, with the effect of boosting each component to enable more powerful tracking capabilities in the presence of disappearing targets. The overall re-id system is meant to be end-to-end (detection performed on context images, feature extraction, re-identification) according to the unconstrained real world scenario.

Outcomes

Over the last 12 months our work on wide-area pattern-of-life learning and anomaly detection has continued in two-directions.

Firstly, we have undertaken some initial work on unifying pattern-of-life learning, anomaly detection and tracking. This work focused on familiarisation with the Hypothesised filter for Independent Stochastic Populations (HISP) [R10] and with fully probabilistic methods for pattern-of-life modelling. The combination of a modified HISP filter with online Kernel Density Estimation [R11] (oKDE) has shown that a compact probability density function can be constructed to model spatial context in an online setting. In contrast to standard KDE, which requires all datapoints to be maintained indefinitely, oKDE has shown that it can model the true spatial distribution of target pattern-of-life with only 12% error while compressing the data that must be retained by 97%.

Secondly, we have been further developing and evaluating our Gaussian Chains algorithm for wide area pattern-of-life learning and anomaly detection. More specifically, by incorporating elements of the oKDE algorithm we have been able to eliminate several free-parameters from our pattern-of-life learning algorithm. Furthermore, we have now been using Conformal Anomaly Detection (CAD) by updating the Sequential Hausdorff Nearest Neighbour (SHNN) algorithm [R12] to use on our pattern-of-life learning model. Evaluation has shown that our new algorithm (GC-SHNN) is able to reduce the runtime of the SHNN baseline by approximately 50% without impacting recognition performance.

The Temporal Anomaly Detection (TAD) enabling contract, which ran from November 2015 to August 2016, saw the team address the problem of anomalous device identification in noisy broadband environments. This work showed that the Symbol Aggregate approximation (SAX) algorithm [R12] performs poorly in such environments, with weak anomalous devices remaining undetectable. Through this project we have delivered four significant outputs: 1) A novel algorithm for removing noise from spectrogram data. 2) Frequency-SAX: An extension of the SAX algorithm to incorporate frequency-based features that permits the detection of transient unknown devices. 3) SEGM: A new image segmentation based approach which performs automatic clustering of spectrogram images enabling anomalous device detection and re-identification. 4) Implementation code for each algorithm has been supplied to Dstl.

With regards to our research on people tracking by re-identification we have implemented a baseline framework which; a) performs convenient pre-processing of the two largest benchmark datasets; b) extracts discriminative deep features by a residual learning-based Convolutional Neural Network (CNN); c) implements original performance evaluation protocols; d) displays ranking results and mistaken predictions. This unit represents a module of a larger end-to-end framework which is being implemented in collaboration with Dr Yang Hua, a researcher from Queen's University, Belfast.

Progress

Unifying pattern-of-life learning and Tracking

We have commenced our work on unifying target tracking, pattern-of-life (POL) learning, and anomaly detection into a single framework. Our progress thus far has focused on finding an effective technique for modelling POL as a probability density function that can be directly constructed and re-utilised by a target tracking algorithm. For the underlying tracking algorithm we are currently using the Hypothesised filter for Independent Stochastic Populations (HISP) [R10] developed in-part by WP5.

Initial experiments have shown that online kernel density estimation (oKDE) [R11] is able to model POL from a simulated traffic intersection with low error. Figure 13 and 14 shows a simulated scenario with 60 tracks (black) under noise (red). oKDE is able to model the POL in this scene using 71 Gaussian components; a 97% reduction in the number of datapoints over the standard KDE algorithm. The Jenson-Shannon distance metric from the full KDE shows an error of 12% is introduced by the oKDE algorithm.

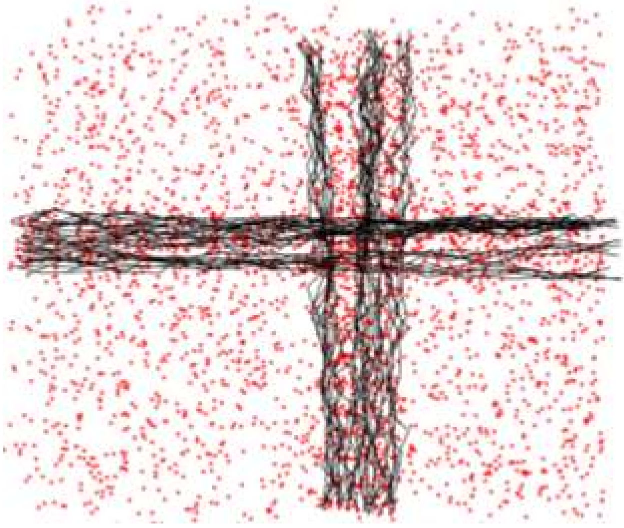

Figure 13: Simulated tacks at a traffic intersection

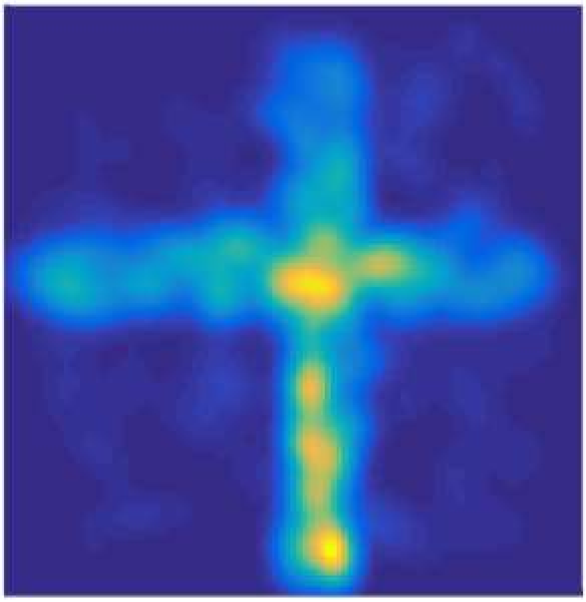

Figure 14: The online KDE probability density function.

Pattern-of-life learning and Anomaly Detection

Prior work on pattern-of-life (POL) learning and anomaly detection has largely focused on easy applications such as street intersections [R13] and transport hubs [R13]. However, long-range wide-field sensors are increasingly being deployed (e.g. wide area motion imagery - WAMI). Processing such data raises several `big-data' challenges that have not been solved. By `big data' we mean streaming sensors/sensor networks that monitor very wide areas (≥ 35^2 km), having many targets (≥ 1000) per sample (e.g. video frame), and monitoring over long durations (days/weeks/years).

Our parallel work on POL learning for unification with target tracking led us to identify an integration opportunity with our Gaussian Chains algorithm previously reported. Our new approach uses a spatial clustering step based on the oKDE algorithm to ‘down-sample’ trajectories. This is achieved by using online Gaussian Mixture Model learning, which processes the data in a bottom-up manner and without the need to specify the number of mixture components.

Figure 15 illustrates the overall process in which three trajectories (comprised of 168 positions) are down-sampled to just 3 mixture components. The number of point-wise distance calculations that must be performed at the anomaly detection stage (not shown) is reduced from 168^2 to 11^2 (a 233% reduction).

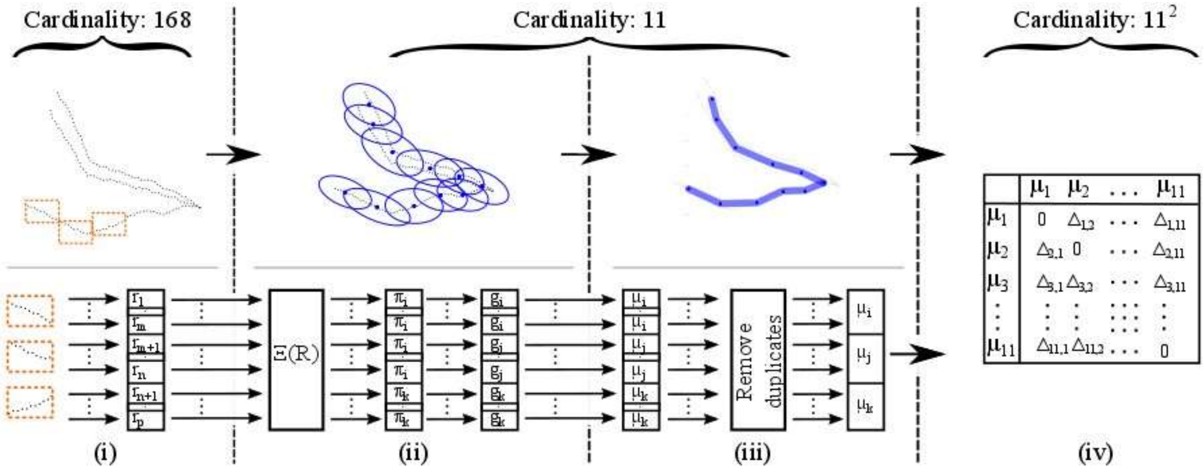

Figure 15: Reducing the number of trajectory positions through cluster. i) Top: Three input trajectories. Bottom: Three sets of positions. ii) Top: Partitioning of all trajectory points into 11 mixtures. Bottom: Mapping of input positions to mixtures. iii) Top: New trajectory representation based on μ of each mixture. Bottom: Removal of contiguous duplicates in output representation. iv) Distance matrix calculated between each mixture.

We have evaluated the performance of our pattern-of-life model by integrating it into the Sequential Hausdorff Nearest Neighbour (SHNN) Conformal Anomaly Detector, henceforth referred to as the GC-SHNN. As a comparative baseline we used the standard SHNN algorithm, which operates on the raw trajectory input rather than the down-sampled trajectories. We report the improvement in anomaly detection F_1-score, defined as:

$$F1 = 2 \cdot \frac{PPV \cdot TPR}{PPV + TPR}$$

The F_1-score is an evenly weighted combination of positive prediction value/precision (PPV) and true positive rate/sensitivity (TPR) and thus is a useful metric for comparing overall classification performance.

Benchmark comparison

In Figure 16 we report mean improvement metrics on three public datasets: Piciarelli (PIC), LAB, and the Laxhammar (LAX). In each experiment only three trajectories were used for initial POL training, with the remaining trajectories used for testing with online adaptation. It is clear that for all three datasets our algorithm delivered substantial improvements in runtime with reductions of 27%, 80%, and 52%. All improvements are significant at 95% confidence. With respect to changes in F_1-score, the only notable difference was on the Piciarelli dataset, however, it was not statistically significant at 95% confidence.

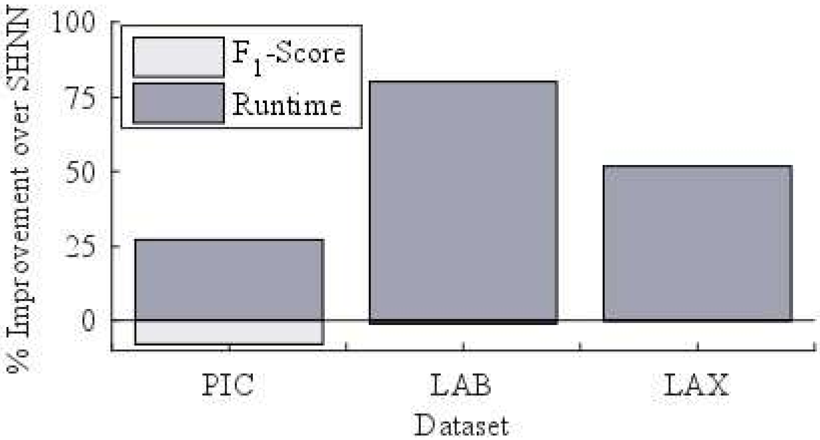

Figure 16: Reducing the number of trajectory positions through c

Effect of trajectory length

To explore the effect of trajectory length on algorithm performance we generated a new dataset using the same generator that produced the Laxhammar and Piciarelli datasets. We used 10 trajectory classes to be consistent with Laxhammar, and used a subset size of 1000 trajectories. For each subset we randomly sampled and replaced ~1% of the normal trajectories with random trajectories. One-hundred subsets were generated per trajectory length, and trajectory lengths of 16,26, 66 were used. As with the benchmark experiments, each algorithm was evaluated by using the first three trajectories as training data with the remainder used for testing with online adaptation.

Mean performance metrics can be seen in Figure 17. It is clear that SHNN performs worst in terms of runtime:trajectory length, increasing x16 between trajectory length 16 to 66. Our approach shows a gradual increase in runtime. In terms of F_1-score SHNN performs consistently at ~0.7 while our algorithm suffers slight degradation on shorter trajectories (~0.65) while reaching the same performance of ~0.7 by trajectory length 36.

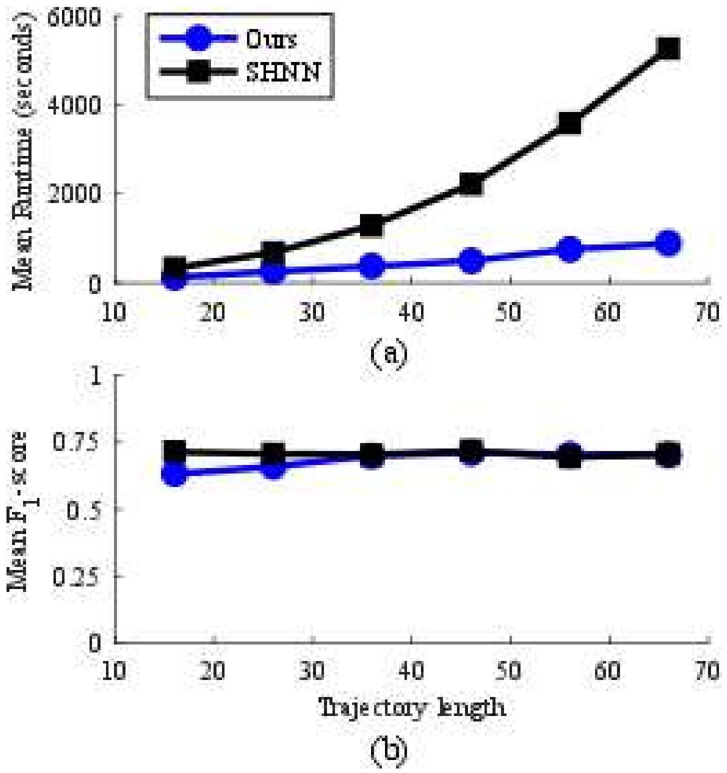

Figure 17: Average a) F_1-score and b) runtime metrics of each algorithm as a function of trajectory length.

Wide area GPS data

To address limitations with previously available datasets, a dataset of 766 GPS trajectories was gathered by ourselves over a seven month period and henceforth will be referred to as the HW-POL1 dataset. Trajectories represent the movements of a single individual going about their daily routine and were gathered at a sampling frequency of 5 seconds. Each trajectory was hand-annotated with a ground truth label indicating if it was spatially and/or temporally anomalous which, although subjective, provides ground truth that can be compared to.

A spatial anomaly was defined to be any trajectory/subtrajectory that the collector perceived to be travelled less than once every 6 months, and a temporal anomaly was any trajectory travelled at a time of day that was perceived as inconsistent from the person's historical behaviour. Figure 18 illustrates two concrete examples from the dataset: Area (A) shows a trajectory labelled `normal', and its corresponding frequency and time-of-day distribution. Area (B) illustrates a spatio-temporal anomaly for a trajectory only travelled once. Note that by definition a spatial anomaly must also be a temporal anomaly, however, a temporal anomaly need not be spatial. For example, traversing the trajectory in (A) at 2am would only be a temporal anomaly. In the experiments that follow we consider spatio-temporal anomalies only while encouraging the use of the temporal information for future work.

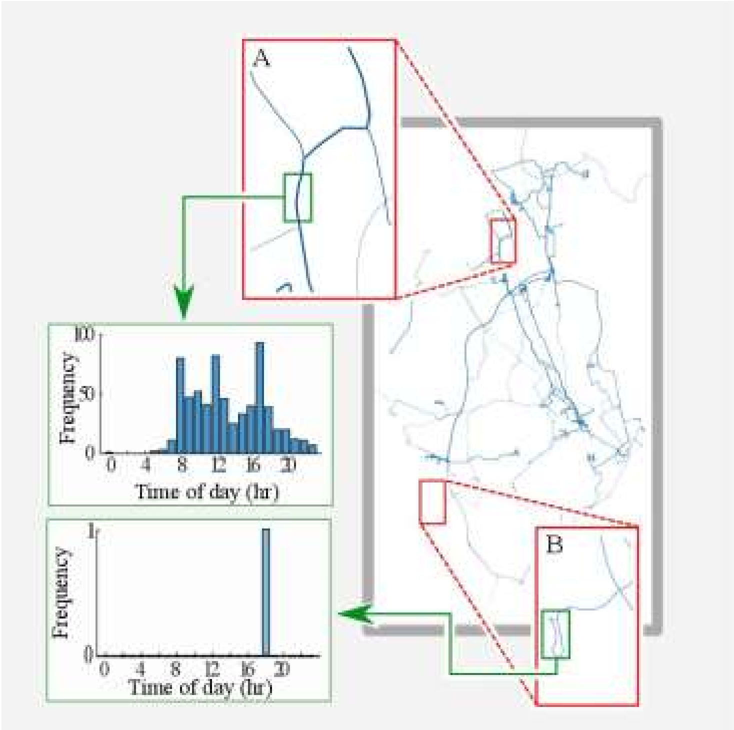

Figure 18: Example trajectories in the HW-POL1 dataset. Region (A) highlights a trajectory labelled as normal and having a wide temporal distribution. Region (B) highlights a trajectory only observed once and labelled as anomalous.

Table 1 compares the performance of our algorithm against the SHNN baseline. Our algorithm outperformed the SHNN in all metrics, although particular attention is brought to the F_1-score and runtime. We achieved an F_1-score of 0.72 compared to 0.65 for the SHNN (11% improvement), and a runtime of 150 minutes compared to 404 minutes for the SHNN (a 270% improvement). We speculate that the gain in F_1-score is because our model generalises better to new data.

Table 1: Anomaly detection performance on the HW-POL1 dataset.

Method	TPR	PPV	Acc	F_1-Score	Runtime
Ours	**0.73**	**0.72**	**0.87**	**0.72**	**150m**
SHNN	0.71	0.61	0.77	0.65	404m

People tracking by re-identifications.

The progress on this piece of research deals mainly with the definition of the feature extraction net, implemented by a residual learning-based net with 50 layers (ResNet50). We have evaluated the performance of the re-identification (re-id) system in terms of ranking according to two approaches: 1) computing the cosine similarity between the extracted feature vectors; 2) learning a metric in the features space by the Joint-Bayesian method.

Our baseline rank 1 re-id rate achieved with the normalized Euclidean distance is 71% (Figure 19) which represents state-of-the-art result considering that no alignment of pedestrians' full body is applied. The same evaluation has been carried out on a more basic CNN, the Caffe net, which perform much worse (as shown in Table 2) because it is significantly shallower than ResNet50. Indeed, since the Caffe net does not exploit the residual learning paradigm it suffers decreasing accuracy as depth increases. We are also performing an analysis of the wrong predicted samples to make the re-id weakness of the current net emerge, in particular with regards to pedestrians' poses (Figure 20).

metric	mAP	rank 1
Caffe net - cosine	14.58%	31.77%
Caffe net - joint-Bayesian	16.33%	32.72%
BoW + Geo + Gauss	14.74%	35.83%
ResNet50 - cosine	45.49%	71.02%

Table 2: Performance of the ranking-based re-id system implemented by a) Caffe net and b) ResNet50, applying different distance metric (comparison with the non-deep features case).

We are going to formalize the current results completed by the outcomes of the re-id Siamese net in a joint paper with Yang Hua to be submitted to ICCV/BMVC.

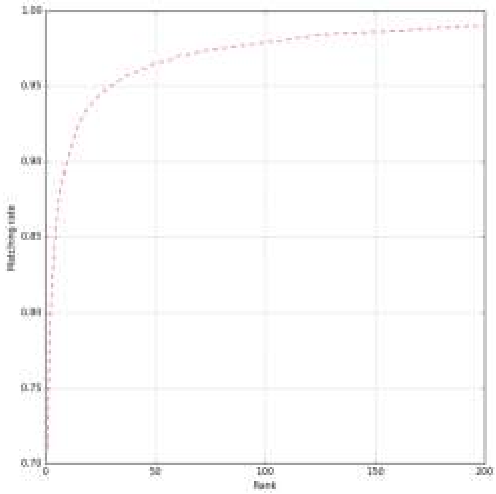

Figure 19: CMC for people re-id using ResNet50 and cosine distance on Market-1501.

Figure 20: Top ranked images corresponding to their probes (results for the rank 5 case).

Future Direction

In the short term we will be continuing to evaluate our algorithm on a second baseline algorithm: Discords. This work has already commenced with the algorithm implemented and several sets of results obtained. Initial results suggest that substantial improvements in both F1-Score and runtime are achieved by our algorithm, however, full metrics for all datasets are not yet available. Furthermore, we will be evaluating the performance of all three algorithms on the Wright Patterson Air Force Base WAMI dataset.

For our work on target tracking with re-identification we will be implementing the remaining modules of the overall re-id pipeline. Specifically, the detection net and the re-id nets. Our aim is to overcome the limitations of the solution proposed in the end-to-end framework in [R14]. The re-id net is meant to be a Siamese cnn adapted to the re-id problem used for learning a metric in the features space. We are considering adapting the fully convolutional net in [R17] (but used for tracking in the paper). The detection net should be a light-weight network specific for pedestrians like MTCNN [R16] for face detection, in order to allow for the application of a deeper structure to the subsequent feature extraction net (for gaining more discriminating power). The novelty of our work will be in concatenating and learning jointly the detector and the Siamese net, introducing a convenient loss function. The ultimate goal is to fit the re-id pipeline into the framework for estimating the targets reappearance distributions in the network of cameras (Figure 21). We are planning to present a seminar at Roke Manor in May about the conference paper we are going to submit to ICCV/BMVC in the next few months.

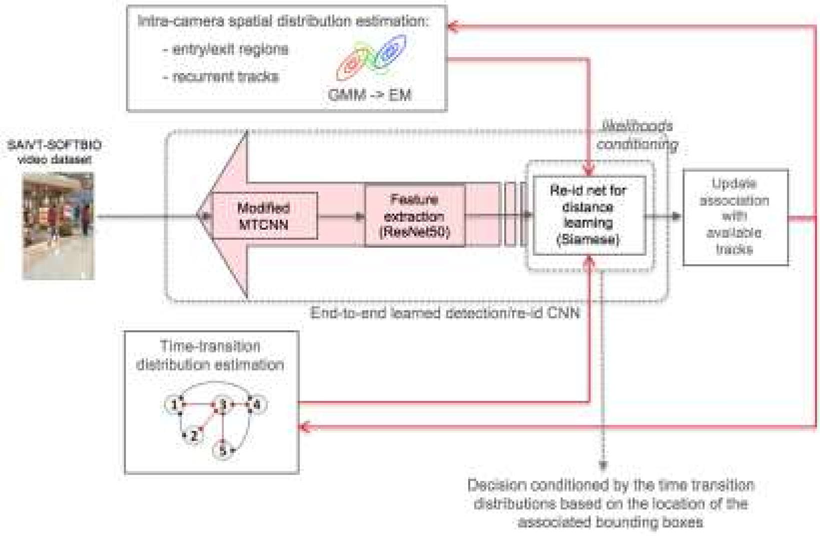

Figure 21: Integration of the end-to-end re-id system in the time-transition distribution estimation framework.

WP5 Estimation Framework for Multi-target Detection/Tracking and Sensor Management

Research Leader: Daniel Clark

Academics: Yvan Petillot, Mike Davies

Research Associate: Emmanuel Delande

PhD Students: Jose Franco, Alexey Narykov, David Cormack

The aim of this work package is to provide a unified framework for multi-object Bayesian filtering (i.e. multi-target detection and tracking), multi-sensor data fusion and sensor management. Various domains of applications motivate the developments in this work package, with a main focus on the domain of Space Situational Awareness (SSA), a challenging environment that has been the topic of a growing interest from the defence communities in the UK and abroad.

This work package articulates around both theoretical and practical developments. We propose significant contributions to the development of a new estimation framework for stochastic populations [R19], mostly focused on the design and implementation of the DISP filter [R41] and the HISP filter [P48] and information-theoretic tools for sensor management [R18], as well as various contributions [P42, R20, P43, P44, R21, P46] on the enrichment of the well-established FISST framework [R22]. Applications on these solutions, including the exploitation of real data, is most notably focused on the SSA domain [P45, R23, R24, R25, P48]. An ongoing collaboration with the biology department of Heriot-Watt exploits some of the results of this work package for non-defence applications [P49].

Outcomes

The first outcome is the explicit connection between a) our early work on regional statistics and information-theoretic gain for sensor policies, and b) the tracking algorithms derived from the estimation framework for stochastic populations, articulated into a single comprehensive paper targeting a journal on automatic control. A submission to IEEE Transactions on Automatic Control has been proposed.

The second outcome is the implementation of the new DISP filter in native C++ with standardized I/O structures in XML format, and the design of a testing bench in MATLAB using the same standard. Following this, relevant target detection/tracking scenarios can be implemented in MATLAB, processed by filters once their I/O structures are encapsulated in the same XML format, and finally displayed by MATLAB for testing purposes.

The third outcome is the implementation of an approximated version for the DISP filter, the HISP filter, adapted to large scale scenarios where the number of targets precludes the exploitation of combinatorial solutions such as the DISP filter. It has been successfully adapted to the context of SSA, where the number of satellites and orbital debris is one of the main challenges to be addressed for the construction of a catalogue of orbiting objects.

Progress

Most of the achievements during the previous years had focussed on theoretical developments proposing new capabilities for tracking algorithms and the extraction of statistical information for the assessment of sensor actions. A more pedagogical approach followed during this year, and is to be continued next year, in order to increase the visibility of our filtering solutions and illustrate their application to relevant scenarios.

The older version of the DISP filter submitted for SSA applications [P45] has been generalized in order to propagate information on yet-to-be-detected targets. This new version of the DISP filter allows for the probabilistic description of sensor actions focussing either on a) the exploitation of previously-detected tracks, or b) the exploration of the surveillance area for yet-to-be-detected targets, and thus facilitates their comparison and integration in a unified sensor management policy. We have articulated the connection of this new filtering solution with our previous work on regional statistics for set-based filters [R20, R21] and information-theoretic sensor management for stochastic populations [R18]. A comprehensive description of the DISP filter in its more general form, with explicit connections with a) approximated solutions as exploited in [P45], and b) statistical tools as in [R18, R20, R21], is now secured and was submitted to the IEEE Transactions on Automatic Control [P41]. The submission was unsuccessful, however, and the reviews suggest that the highly theoretical content of the submission did not fit the remit of IEEE journals. Another strategy is proposed for the coming year (see next section for more information).

The original implementation of the DISP had to be revisited for the more general version presented in [P41], and we saw this as an opportunity to a) Move away from the original C-code embedded in Matlab in order to increase the computational efficiency of the algorithm, and b) Format the input/output structures of the algorithm, in order to facilitate the comparison with other algorithms (including solutions external to this work package) and facilitate the implementation of future algorithms. As a consequence, the new DISP filter is now available in native C++ with I/O formatted in XML.

In order to handle large scale SSA scenarios involving hundreds or thousands of targets, we have implemented an approximated version of the DISP filter, the HISP filter [R19], and presented our results in the 27th Space Flight Mechanics Meeting [P48]. Renewing the successful experience of the past two years, Emmanuel Delande visited Carolin Frueh's research team in Purdue University (U.S.A.) for two weeks in February 2017. The two main outcomes of this visit are as follows. First, the most important physical perturbations were incorporated to the orbital prediction model in order to improve the accuracy of Bayesian tracking filters in the context of SSA (applicable to the HISP/DISP filters, but also to track-based approaches such as the MHT filter, or set-based approaches such as the PHD filter). Second, a collaboration has been initiated in order to exploit the HISP in a sensor management scenario for the surveillance of the satellites on the geosynchronous (GEO) belt with a very limited number of telescopes.

From our collaboration with Isabel Schlangen, a PhD student, we have submitted a journal paper to IEEE Transactions on Signal Processing presenting a new version of the Probability Hypothesis Density (PHD) filter [R22] with Panjer-distributed prior information [P42], and exploited our previous works on the joint estimation of sensor state and multi-target state (see for example [R24]) for the estimation of the sensor clutter rate in a non-defence application related to molecular biology [P49].

We have also compiled new lecture notes [P43] for tutorial sessions delivered this summer at the 2016 UDRC Summer School in Edinburgh, and the 2016 Fusion conference in Heidelberg (Germany), and prepared a paper on the processing of data from the Space Geodesy Facility in Herstmonceux [P47] for the SSPD conference this year. Finally, two journal papers have been accepted for publication, one [P43] for IEEE Transactions on Signal Processing in collaboration with the University of Colorado Boulder, another one [P44] for IEEE Transactions on Aerospace and Electronic Systems.

Future Direction

Following the unsuccessful submission [P41] of the paper presenting the general DISP filter and articulating the connections with our previous work on regional statistics and information-theoretical gain for sensor management, the key priority of the following year is to develop another strategy to communicate our results on that topic. Since the reviews suggested that the highly theoretical content of our submission was out of scope for the IEEE Transactions on Automatic Control, we propose a clearer separation of the contributions through two submissions. A first paper will focus on the more theoretical aspects of our work and highlight the advantages of the estimation framework for stochastic populations for sensor management scenarios: to the best of our knowledge, it provides the first unified description of previously-detected and yet-to-be-detected targets in a unified and coherent estimation framework, and the DISP filter is directly relevant to sensor management problems in multi-target tracking scenarios where exploitation tasks must be weighted against exploration tasks. A second paper will focus on the practical description of the DISP filter, provide a pseudo-code of the algorithm, and illustrate its application on a simulation-based sensor management scenario where the threat of incoming potential targets is translated into the probabilistic modelling of the population of yet-to-be-detected targets. The first paper will be submitted to the SIAM Journal on Control and Optimization (SICON), while the second will be submitted to IEEE Transactions on Automatic Control or IEEE Transactions on Signal Processing.

Due to its linear complexity in the number of maintained tracks and the number of collected observations, the HISP filter seems particularly adapted to the context of SSA, where the combinatorial solutions such as the MHT and DISP prove to be intractable on realistic scenarios. Our early results on the exploitation of the HISP filter for SSA scenarios are very promising, and we intend to integrate the HISP filter to the C++/HTML framework we developed for the DISP filter in order to test the HISP on a relevant SSA scenario involving hundreds of targets. We will rely on the expertise of Carolin Frueh's research group in Purdue University for the generation of realistic data representing orbital trajectories and sensor models; we will also further collaborate with her team and propose our expertise in target tracking algorithms to approach the sensor management problem they formulated for the surveillance of the geosynchronous belt with telescopes.

WP6 Efficient Computation of Complex Signals Processing Algorithms

Research Leader: John Thompson

Academics: Andrew Wallace, Neil Robertson, James Hopgood,

Research Associates: Cris Rusu, Paulo Garcia, Christian Nunez Alvarez

PhD Student: Saurav Sthapit

The overall aim of this work package 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.

The aim of WP6.3 is to propose algorithms for the scheduling and analysis of a heterogeneous sensor network that operates over time with two constraints: maximize (or balance) the accuracy of the task performed by the sensor network while also ensuring a balanced energy consumption among the sensors (for example, the same sensor is not active for a long period and therefore its power supply is not at risk of being depleted).

Outcomes

In our initial work we consider a sensor network that needs to operate over a fixed number of time instances. At each time instance the goal is to perform a measuring task with a minimum prescribed accuracy while the overall goal is to reduce the energy consumption of the network. We accomplished this task by reducing the amount of times a particular sensor is activated. With this constraint we can balance the usage of the network and guarantee that the majority of the sensors will be operable for the whole lifetime of the network. We formulated the sensor management problem as a binary optimization problem which we solve by convex relaxation methods. The method we propose uses ell infinity regularized reweighted L1 penalty functions in an iterative fashion to create an activation schedule for the whole sensor network over all time instances. The method we propose is computationally efficient to solve (polynomial complexity)[P53].

Progress

A successful collaboration with Andy Wallace and Paulo Garcia from Heriot Watt University on the topic of sensor and hardware design of operating models that balance between the energy consumption of the devices and their computational accuracy. Paulo Garcia proposed methods for generating FPGA memory architectures from both Hardware Description Languages and High Level Synthesis designs that minimize memory usage and power consumption. Based on a formalization of on-chip memory configuration options and a power model, we demonstrate partitioning algorithms that outperform traditional strategies, tacking signal/image processing algorithms such as MeanShift Tracking and Optical Flow; power results are depicted in Figures 22 and 23. This research resulted in a journal submission[P59]: "Optimal Memory Allocation and Power Minimization for FPGA-Based Image Processing" IEEE CSVT. Collaboration with the Rathlin project at Heriot-Watt also resulted in one more journal [PP60] and two conference submissions [P61, P62]: RIPL: “A Parallel Image Processing Language for FPGAs" ACM TRETS and “Enabling Parallel Dataflow Graph Transformations using Petri Net Analysis" ICGT.

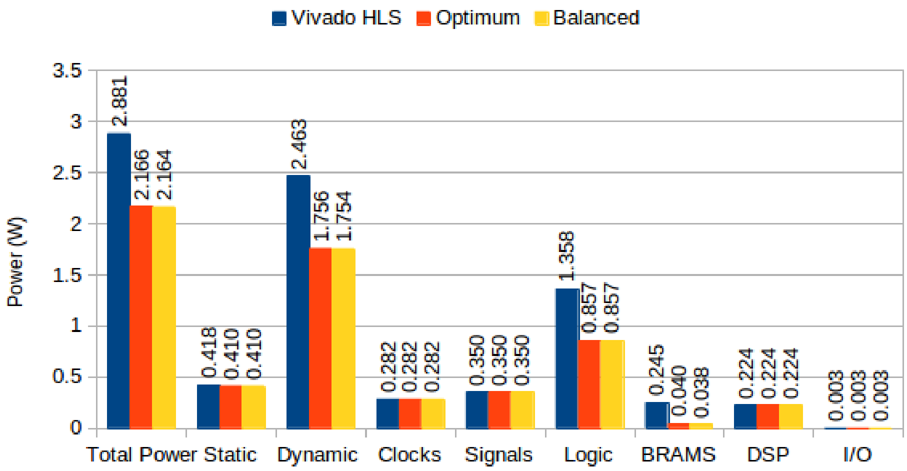

Figure 22: Optical Flow Power consumption on Virtex 7 FPGA

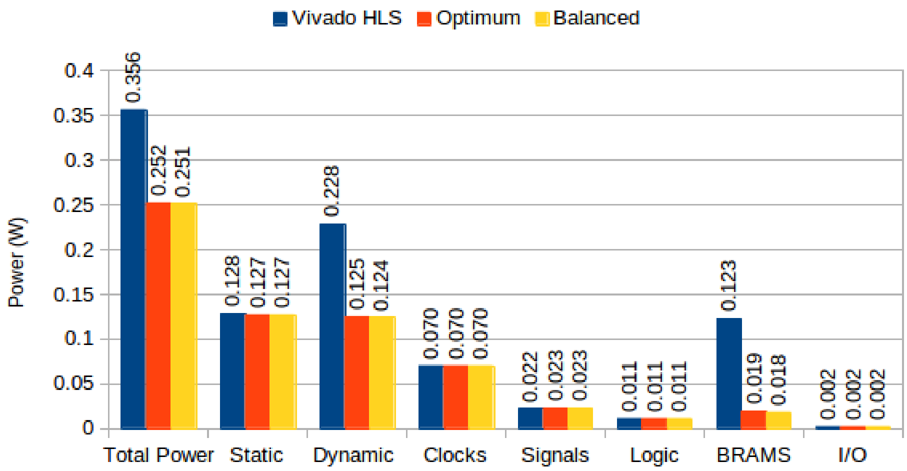

Figure 23: MeanShift Power consumption on Zynq FPGA

"An Interactive Visual Dataflow Tranformation Framework" SCOPESA collaboration with Neil Robertson is underway on the broad topic of sensor management. At the beginning of April both Cris Rusu and John Thompson will visit Neil and his research group at Queen's University Belfast to further explore common research topics and future collaborations.

Cris Rusu has been supporting the activities of the MASNET project by attending meetings and working with the research team. This project has developed a simulation of radio frequency scanning using a network of sensors using statistical channel models. The project is now moving towards more realistic modelling of the radio environment using a ray tracing package and performing experimental evaluations using National Instruments radios.

The research activities have materialized in the following publications:
3 journal papers submitted [Learning Fast Sparsifying Transforms"[P56], "Sensor management with time, energy and communication constraints [P55], Power Reduction on Image Processing on FPGAs [P57] has been submitted to IEEE CSVT"].

One journal paper “Approximate Factorizations of Matrices with a View to Computational Efficiency” is under preparation.

The following conference papers were submitted to EUSIPCO 2017: “On the use of tight frames for optimal sensor placement in time-difference of arrival localization”[P58] and “Learning Fast Sparsifying Overcomplete Dictionaries”[P59]. The conference paper “Learning Fast Orthonormal Sparsifying Transforms” was submitted to SPARS 2017 [P34].

Future Direction

We are planning two research directions, to take the current work further.

We are proposing new, numerically efficient ways, of constructing matrix factorizations in signal processing and machine learning applications. Our goal is to build algorithms that are applicable in O(n log n) memory and operations.

Regarding the sensor management work, we are proposing new algorithms to deal with surveillance, target tracking and mission planning tasks (especially when concerned with UAV and ground units).

With regard to the work on computational off-loading in sensor networks, Saurav is continuing his PhD work on workload balancing in battery powered nodes. After the EUSIPCO paper, he is working towards writing a journal paper. We are considering network of queues and various central and distributed algorithms.

The work on trade-off between energy consumption and accuracy of hardware operations will be developed further. We are currently working on approximate computing on programmable processors, exploring opportunities to trade accuracy for power efficiency.

Management

This year we have built successfully on our research collaborations. Murat Uney from WP2 has travelled to the Sensor Informatics and Biomedical Technology Laboratory of Aalto University, Finland and explored collaboration opportunities in spatio-temporal signal processing with them. In WP3 Yan Pailhas has worked closely with CMRE and NATO on collecting data on classified objects that can be used for our research. WP3 researchers won the EPSRC grant in collaboration with Newcastle University and the University of York with industry representation from Proserv (Nautronix), Subsea 7 Limited and Technip-Coflexip UK Holdings Limited on acoustic sensor networks (USMART). In WP4 Puneet Chaabra has a placement working at Carbomap Ltd on LiDAR signal processing. In WP5, Daniel Clark started his RAEng secondment with Dstl in Portsdown in February and Emmanuel Delande undertook a 2 week research visit to Purdue University (U.S.A.). In WP6, Cris Rusu is planning a trip to Queen's University Belfast to collaborate his research on sensor management.

We are now into year 4 and we have actively engaged with the following 18 companies in working together on taking research forward.

We commenced 8 related industrial case studentships working on signal processing for defence topics and there are 4 more PhD studentships in discussion. In addition to this, we are in discussion with Dstl for a possible PhD studentship to take forward the Raman work.The latest three studentships are Fast Lidar imaging systems for cars with ST Microelectronics and Andy Wallace, Non-linearity in the RF Sensing Chain with Leonardo and Bernie Mulgrew and Real-time implementations of tracking and estimation algorithms for radar systems with Leonardo, Daniel Clark and Bernie Mulgrew.

We were successful in the following CDE with Kawasaki on the tracking of underwater pipelines using wide band sonar (WP3) and tracking algorithms (WP5). We intend to put in two applications to the CDE defence accelerator call on joint multi-aspect sensor registration and tracking of small targets and automatic sensor calibration and online verification.

We have spent time ensuring that unclassified data is readily available on the UDRC website especially data that has been used to produce papers. We also have published Dstl data and promoted the Dstl motion imagery data across the UDRC and have had organisations interested in this data citing the following research interests: real time dynamic contrast enhancement for daylight and IR, testing the structure of motion algorithms and possible data source for CDE application, testing and benchmarking of several compression algorithms.

Mehrdad Yaghoobi has been promoted to Lecturer and we are currently interviewing for a replacement research associate for WP1. Christian Nunez Alvarez will finish his post at the end of March and we have appointed Dr Loukianos Spyrou, who starts in May, to carry on the MASNET work. We also welcome David Cormack who has started his PhD project in collaboration with Leonardo and is currently looking at Dstl maritime data. Jose Franco from WP5 has submitted his PhD on tracking objects in orbit. Puneet Chaabra from WP3 and Di Wu from WP1 are currently in the writing up stage.

Engagement

Engagement, communication and dissemination are key to the success of the UDRC. The Edinburgh Consortium have been monitoring websites and newsletter activity as well as updating the data area and the publications pages and keeping the Edinburgh Consortium wiki up to date.

Work has been undertaken with the University's web-development team to implement a new publications page (figure 24) and www.mod-udrc.org/publications. This means that each publication is now entered as a database entry which can be queried and filtered directly from the front-end of the website. This will streamline the workflow in terms of adding new publications and will improve user-experience on the website, making it easier for users to find a single publication and a group of related publications. Another advantage to keeping a database record of publications is that this can easily be migrated should the website be migrated to a different provider.

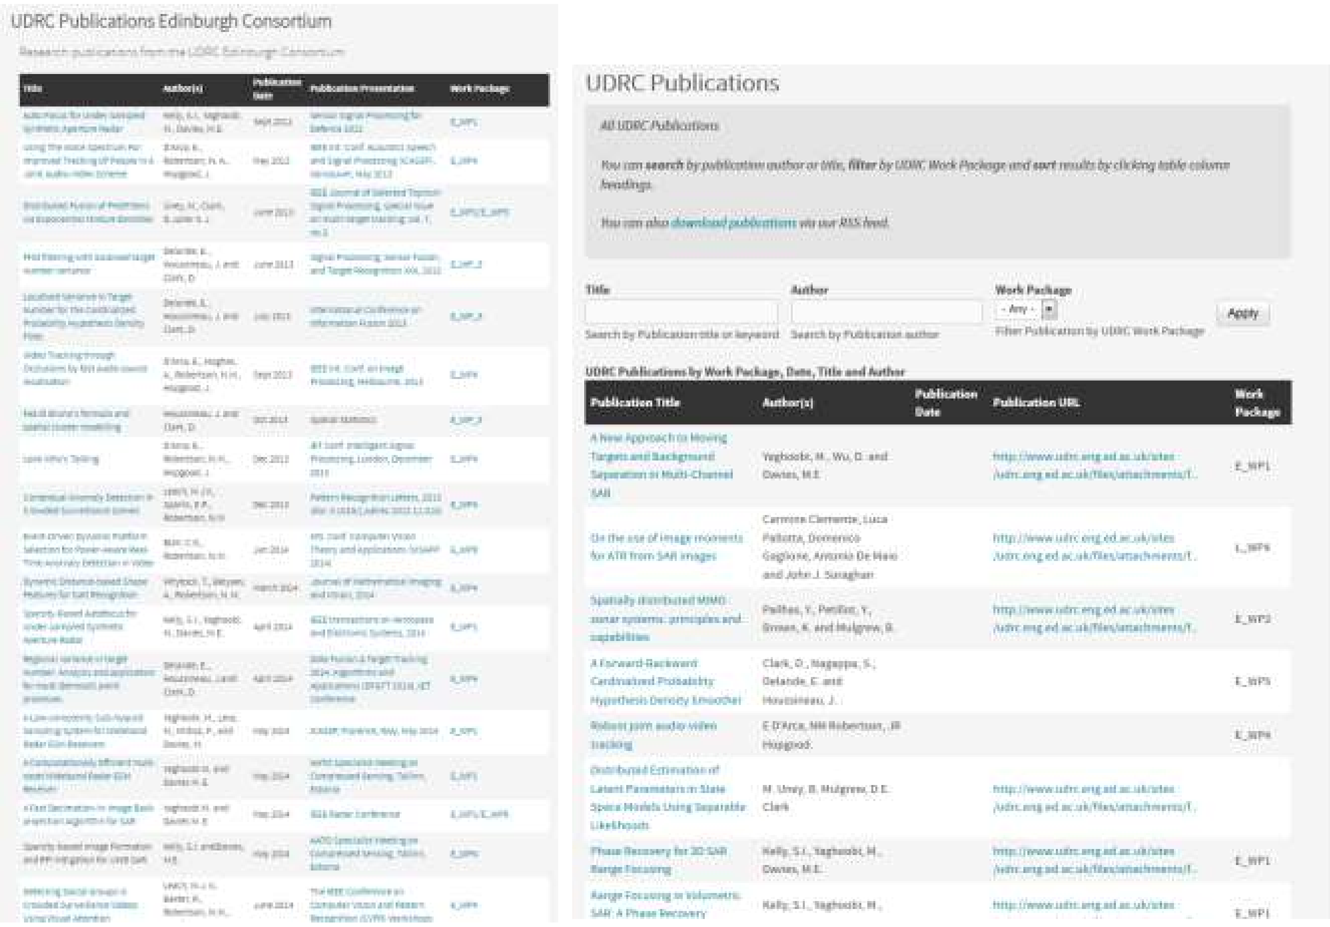

Figure 24: The Publications page before as a long html document (left) and after as single database entries with search form (right).

Websites - Audience and Behaviour

Figures extracted from Google Analytics display a typical seasonality arranged around the two most popular events organised by the UDRC: The Summer School and the SSPD conference (figure 25). These events created peaks in the recorded number of sessions in May (750 sessions recorded on UDRC, 450 on SSPD), June (876 sessions on UDRC, 625 on SSPD) and October (625 recorded sessions on UDRC, 290 on SSPD). A third peak in sessions was recorded in December which can partly be attributed to the December Newsletter.

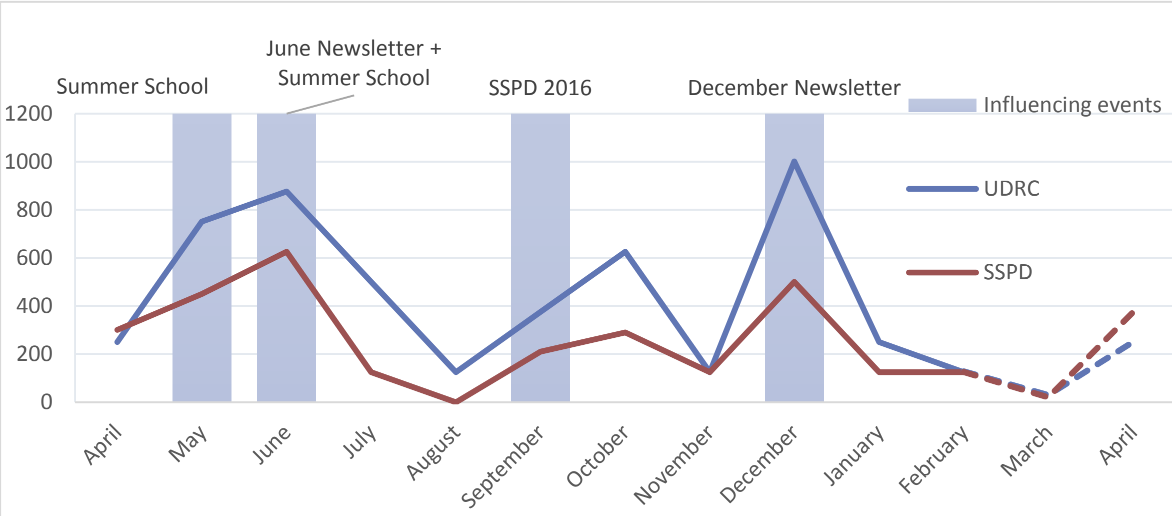

Figure 25: A graphical representation of the number of sessions per month with indications of events that may have influence fluctuations. Source: Google Analytics.

SSPD conference and websites traffic – content analysis

The SSPD conference creates a peak in visits to the UDRC website in the month following the event, possibly as a result of promoting the UDRC's activity and events during the conference. The most visited pages on the UDRC website for the month of October were: People (242 page views), Publication (148 page views) and Events (121 page views). These results appear consistent with the above analysis as attendees to the conference may wish to read further about presenters and their research.

Traffic to the SSPD website also increased after the conference, possibly as a result of attendees looking to download material presented at the conference. A quick analysis of visits by page of the SSPD website for the month of October reveals that the archive pages (26 page views) were amongst the most popular pages for that month along with the programme page (81 page views).

Summer School and websites traffic – content analysis

While the SSPD conference generated traffic after the event, traffic generated by the Summer School was recorded before the event. This behaviour can be interpreted as attendees using the website to find information about the event, accommodation and travel information. This analysis is further supported by the figures extracted from Google Analytics: Most of the visits to the UDRC website happened in early June (418 page views on 8th and 451 on 14th) (figure 24) and were concentrated around the Summer School 2016 Page (435 page views) and the publications (218 page views).

Figure 26: Page views for the UDRC websites for the month of June showing spikes at the beginning of the month.

While it can be expected that the Sumer School generated traffic to the SSPD website, it is not as evident because visits seem more spread over the month (figure 27). However, Google did record high numbers of page views on 29th (73 page views), 30th (81 page views) and 4th July (107 page views) which could have been the result of the distribution of promotional material such as flyers and USB keys to the Summer School attendees.

Figure 27: page views for the SSPD website in the month of June show a peak in visits at the time of the Summer School on 27th to 30th.

Overall figures are very encouraging with a total number of sessions of 5002 for the UDRC websites and 2875 for the SSPD websites. Further, bounce rates (ratio of single interaction visits) of 37.5% (UDRC) and 52.17% (SSPD) are well within acceptable levels for the industry. Low bounce rates and events resulting in visits to the websites, which in turn promote future events is a very positive trend that hints at both the importance and effectiveness of the websites in promoting the UDRC's activities and the SSPD conference.

Returning vs new visitors

This analysis can give some information on how the websites are used. Over the reporting period, both websites have been visited by a majority of New Visitors (figure 28).

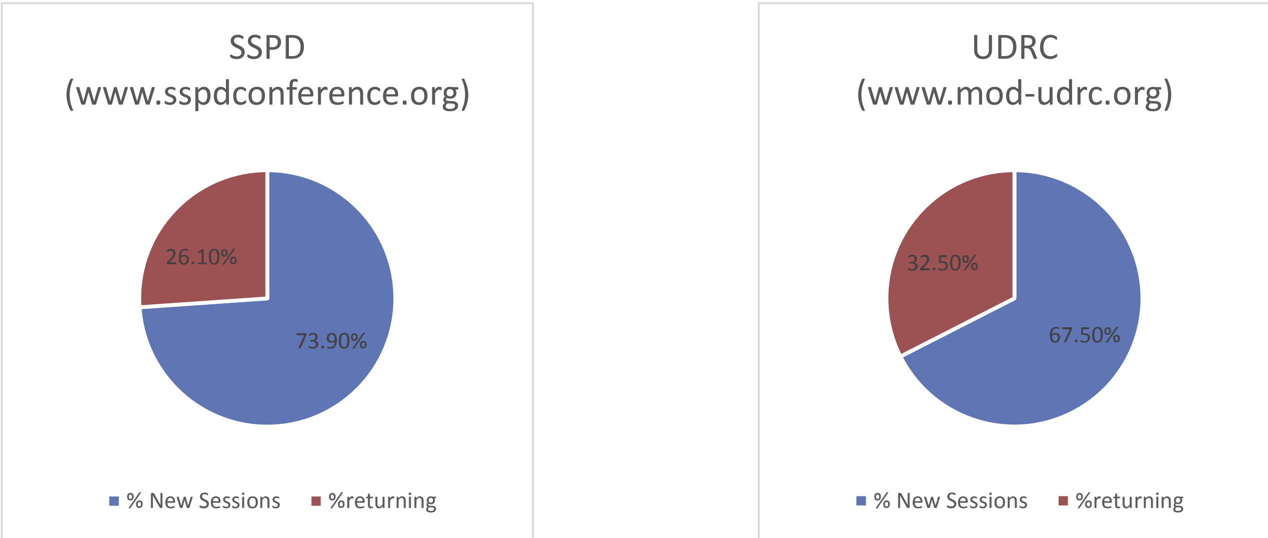

Figure 28: Percentage returning visitors vs new visitors

Considering the highly specialised content of both websites and since these are regularly visited by members of the UDRC, change in this returning to new visits ratio are expected to be largely driven by new visits (figure 29).

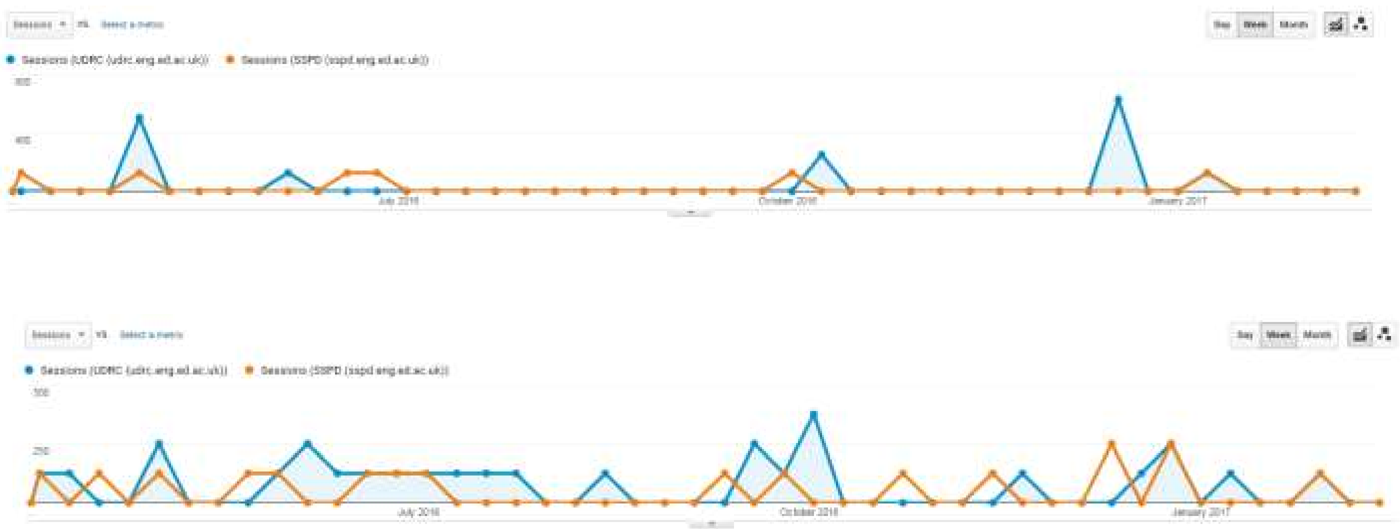

Figure 29: Sessions over the reporting period for returning visitors (top) and new visitors (bottom).

A high proportion of new sessions can be interpreted as a result of online campaigns and appropriate promotion during events thus attracting new visitors to both websites. This is further explained below where the Newsletter, social network and use of the Stakeholders list are discussed.

Acquisition

For both websites, the main acquisition channel was organic search, i.e. typing keywords in a search engine (figure 30). It is therefore important that the websites are appropriately optimised so that they can be found for relevant keyword queries. Figure 31 shows that bounce rates associated to organic search are within acceptable rates which is an indication that, not only do users find the websites via keyword search, but they also interact further with them.

The second most important gateway is “direct traffic”. This is recorded when the URL of the websites is directly inputted into the web browser. Accessing a website via a bookmark also counts as direct traffic. This high proportion of direct traffic could be interpreted as visits by regular users who are likely to have bookmarked or cached URL information in their browser, indicating frequent use. Again, this supports the idea that UDRC members and stakeholders form a core of regular visitors to the websites.

Referrals, or accesses to the website via links from other websites, account for much less traffic in both websites but are nonetheless very important. Figure 31 shows that bounce rates for this type of access are very low, indicating that the majority of users coming from these links found highly relevant content and further interacted with the websites.

Links are an extremely important part of Search Engine Optimisation (SEO). Highly relevant, large website links to SSPD and UDRC websites play a crucial role in the websites “visibility” on the internet. The SSPD website is linked to by highly relevant websites such as www.ieee-aess.org, www.ieee.org and www.spie.org. The UDRC website is also well linked to, with some of the major links being: www.ieee.org, www.rcuk.ac.uk, www.ic.ac.uk, www.surrey.ac.uk.

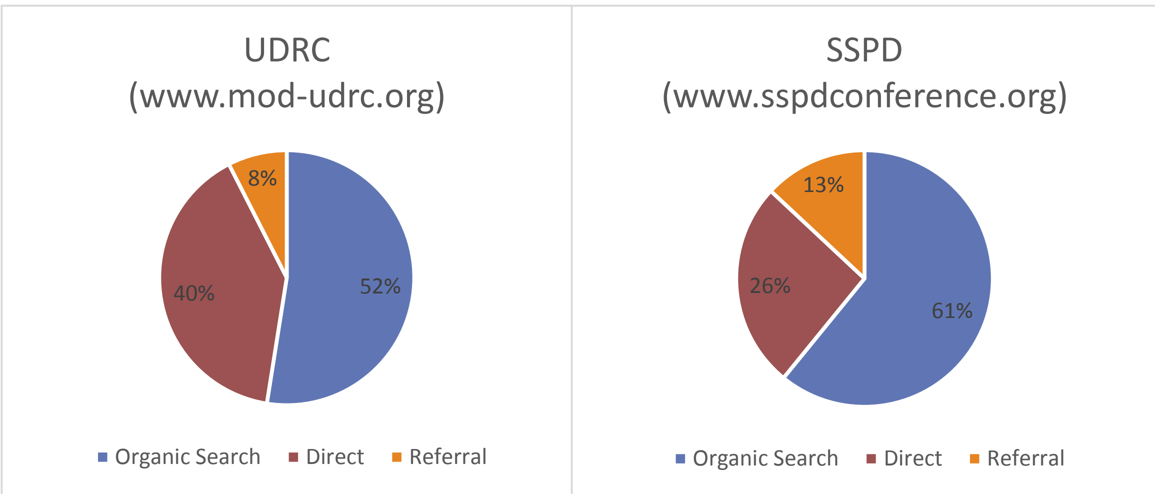

Figure 30: Part of each acquisition channel for each websites.

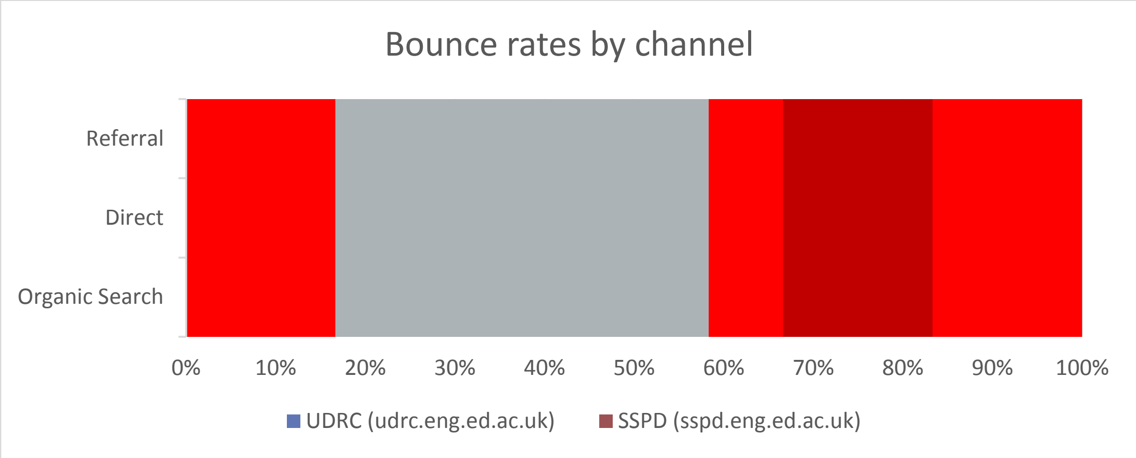

Figure 31: Bounce rates by acquisition channel.

A higher proportion of new sessions to the websites associated with low bounce rates is strong evidence that the websites are correctly optimised to attract interested visitors and presenting them with relevant and engaging content.

Geographic data

Both the SSPD and UDRC websites attract visitors from around the world with visits coming from 21 different countries. Figure 32 shows visits to both websites by location of IP addresses.

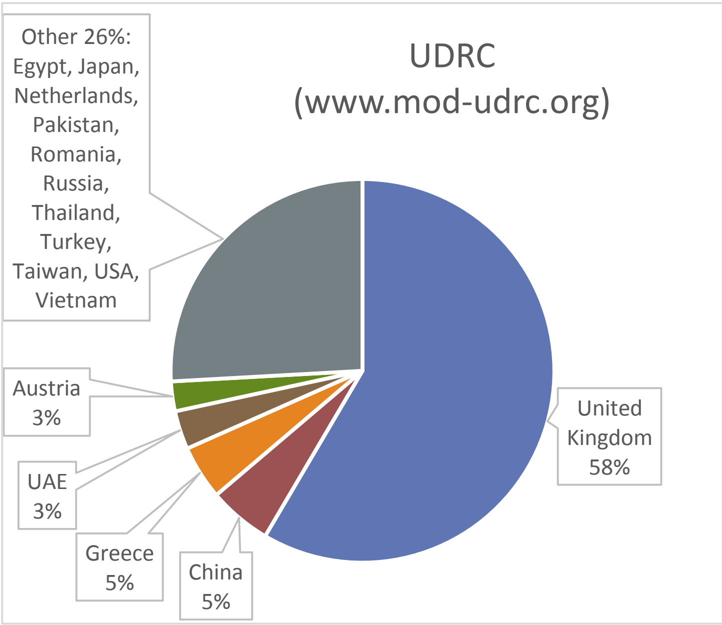

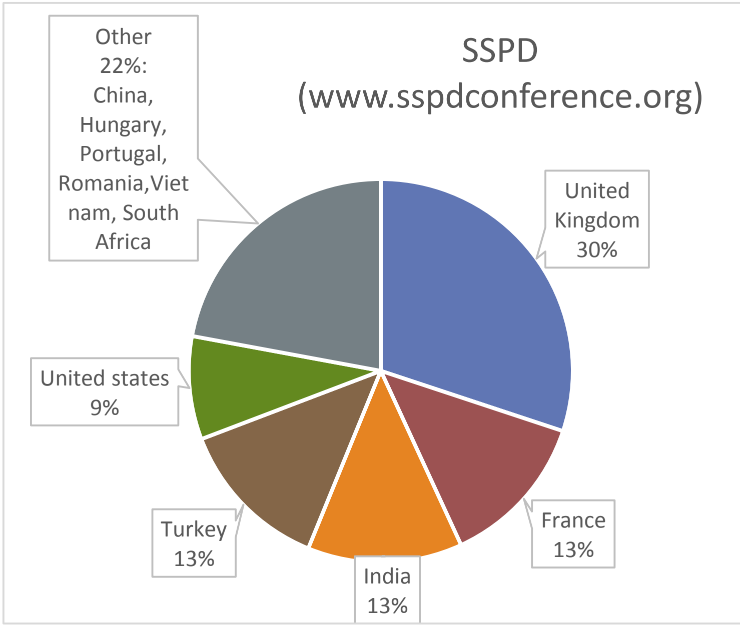

Figure 32: Origin of visiting IP addresses recorded by Google.

The UDRC website has a large majority of visits coming from the UK. This supports the idea of a very specific use of the website by a number of regular visitors likely to be UDRC members or stakeholders. While the UDRC's outreach is international, the core of its activities remains in the UK and this is reflected in the UDRC website figures.

The international aspect of the project becomes all the more evident when looking at figures recorded for the SSPD website. Here, local visits account for a much smaller part while foreign visits account for the vast majority with 70% of visits. Because the SSPD website is used to inform and promote this particular event, local and returning visits are sparser and often associated with announcements made on the UDRC website, newsletters and emails. This is consistent with figure 28 showing a lower proportion of returning visits on the SSPD website (26.10% on SSPD against 32.5% on UDRC).

During the reporting period the websites have recorded a combined 4250 sessions from 20 different countries reflecting the international outreach of the project.

LinkedIn

The UDRC LinkedIn group has experienced a steady growth since its creation in January 2014. During the reporting period an email was circulated to the whole stakeholder mailing list to suggest joining the UDRC LinkedIn group. As a result, membership nearly doubled in the following few days (figure 33). This occurred at the same time as the Summer School 2016 which further helped the sudden increase in membership. Today, the UDRC LinkedIn group counts 176 members. This number can be further increased via promoting the group through events. For instance, a sudden rise in membership was recorded at the same time as the Summer School 2014 (figure 33). Summer Schools appear to be good opportunities to increase LinkedIn membership. Arguably, this is because of the type of audience attracted by the event: Most attendees are younger students who are likely to be more receptive to social media invitations.

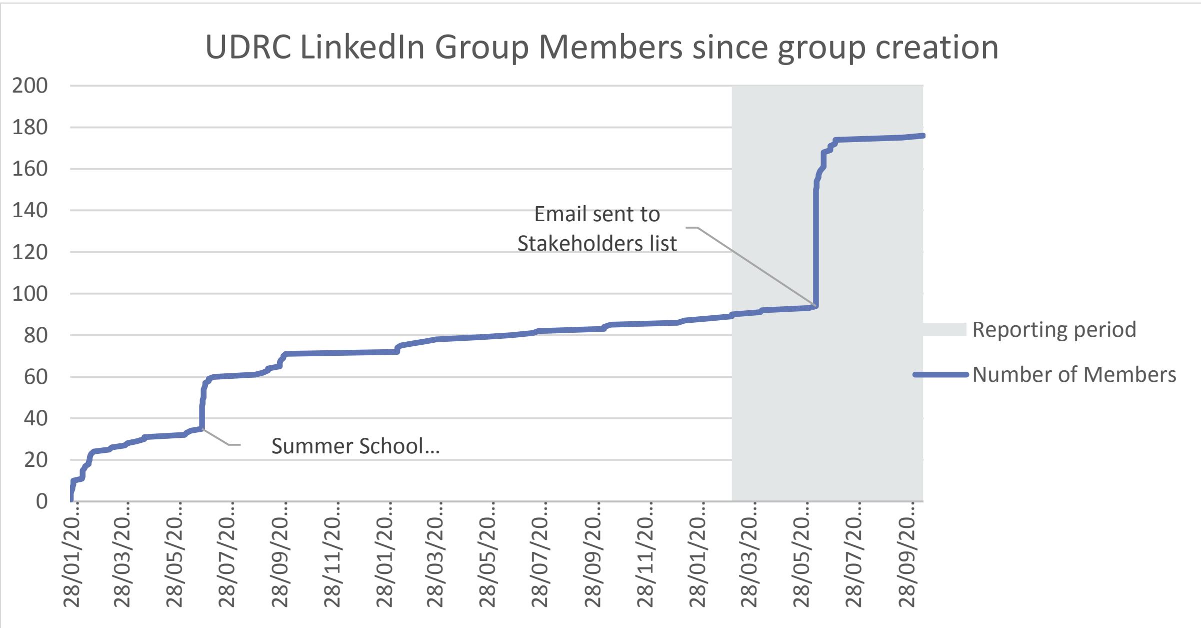

Figure 33: UDRC LinkedIn group membership over time.

Newsletter

The UDRC Newsletter have been playing an important role in keeping the UDRC audience engaged. Figure 34 shows how a newsletter can produce a dramatic increase in visits to the website. The newsletter data presented in figure 34 shows that open rates have been well above industry average for each edition, indicating a strong interest amongst the recipients.

While opening rates appear to have a decreasing trend, the number of opens are increasing over time suggesting that the UDRC Newsletter audience is growing. The decreasing trend of opening and click rates can be explained by the increase in the number of recipients. Indeed, the reporting period has seen the number of recipients grow from 790 to 843, a 6.7% increase.

Furthermore, the last Newsletter, sent on 8th December 2016, appears to have generated 72 direct visits to the Websites. When analysing the Newsletter content, it quickly becomes apparent that most of its content reports on people and achievements. Google Analytics shows that the majority of page views on the UDRC website for the days following the Newsletter were the “People” page (43 page views on 8th December only) and the “Research” page (40 page views on 8th December).

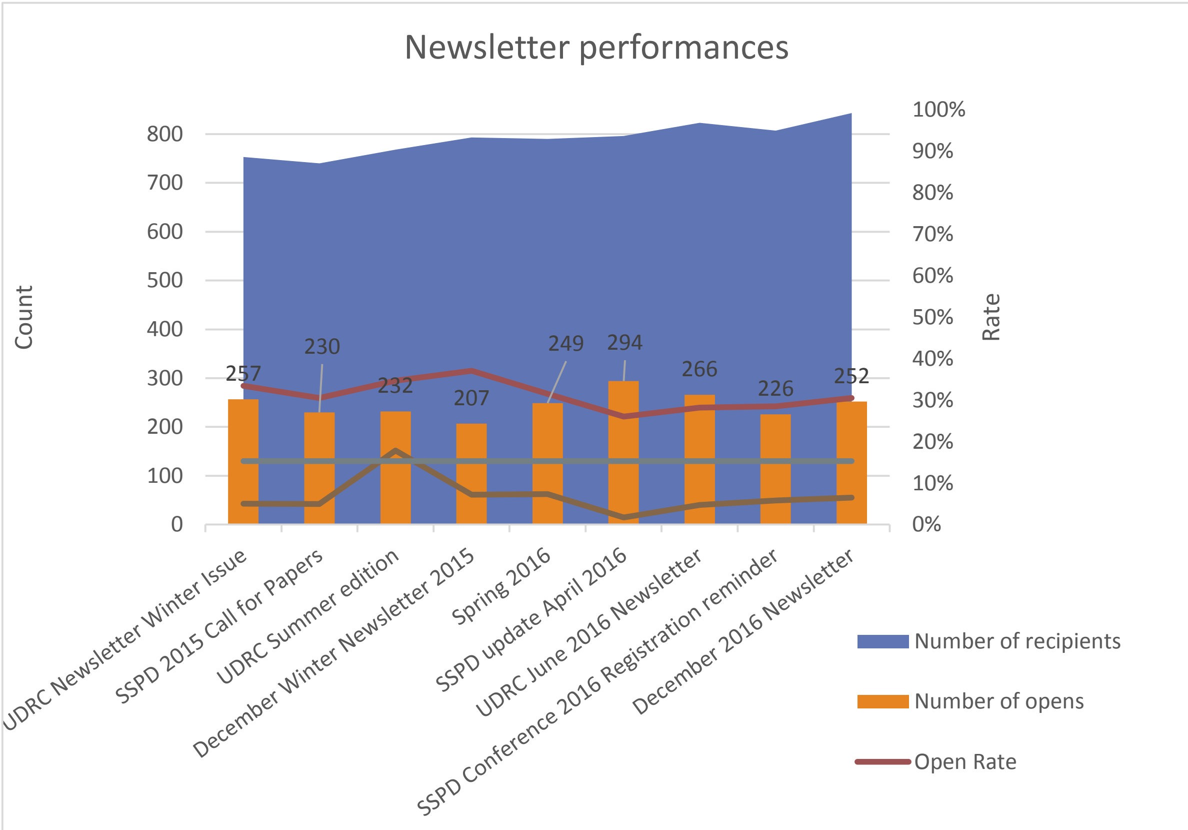

Figure 34: UDRC Newsletter perfornmance looking at open rate and receipients

UDRC wiki

The UDRC Wiki was set up in order to allow all UDRC staff and students to be able to easily access all documents, templates, reports, instructions, minutes, presentations, marketing and publicity documents, used across the Project, in fact, anything that might aid them in their work. The wiki has indeed become a valuable place where all documents can be easily viewed and referred to, and, from the start, it was also anticipated that it would be available as a historical account of the evolution of the administration of the Project (as well as the Research work). It was agreed that any document that was therefore useful in any way would be on the Wiki for all the Edinburgh Consortium to view, share and benefit from, as well as helpful to those applying for the next phase.

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 LSSCN Consortium and Dstl.

Knowledge Transfer Meeting

Our 4th Knowledge Transfer meeting was held at Porton Down in April with over 40 people in attendance. The meeting consisted of a number of workshops where Dstl presented a problem or question and the workshops discussed the problem and possible solutions. Feedback from the day stated that the more successful workshops had questions that were more focused and that these particular workshops generated actions and collaborations.

Workshop titles were:

- Source separation of acoustic data,
- Information fusion for classification from large and heterogeneous data sets,
- Sensor Registration, Spectral source separation,
- MIMO,
- Data analytics problems from the JFIG Centre for Intelligence Innovation,
- EO tracking of extended targets in the maritime environment

UDRC Themed Meeting

Over the past year we have hosted two themed meetings. The first on the topic of Image and Video Processing which attracted 30 people. Dr Toby Breckon from Durham University was invited to discuss his work on real-time target analysis and reporting from infrared imagery within a fused wide-area surveillance context as well as Dr Ian Brown from Thales who discussed the real-time challenges in a military environment. The challenge this year was presented by Roke Manor on the determination of tracks using the WAMI dataset and more specifically determining and labelling parts of each image that are occluded due to buildings. The deadline for this is October.

Our second themed meeting took place on the 23rd November on the topic of Space Surviellance and Tracking at Heriot-Watt University. We had invited speakers Dr Roberto Armellin, Surrey Space Centre, Dr Camilla Colombo, Politecnico di Milano and Waldemar Franczak, Spectator as well as speakers from Edinburgh and LSSSCN.

The next themed meeting will be on the topic of Underwater Sensing, Signal Processing and Communications. This will be held at Newcastle University and the guest speaker will be Dr Yuriy Zakharov from the University of York. This meeting will also introduce the successful collaboration between Heriot-Watt, Newscastle University and the University of York on the USMART grant (Smart dust for large scale underwater wireless sensing).

Summer School

We held another successful UDRC Summer School in June 2016, at the Kings Buildings and 66 students were in attendance across 4 days of lectures. These students were made up of 50% academia, 20% Dstl and 30% industry.

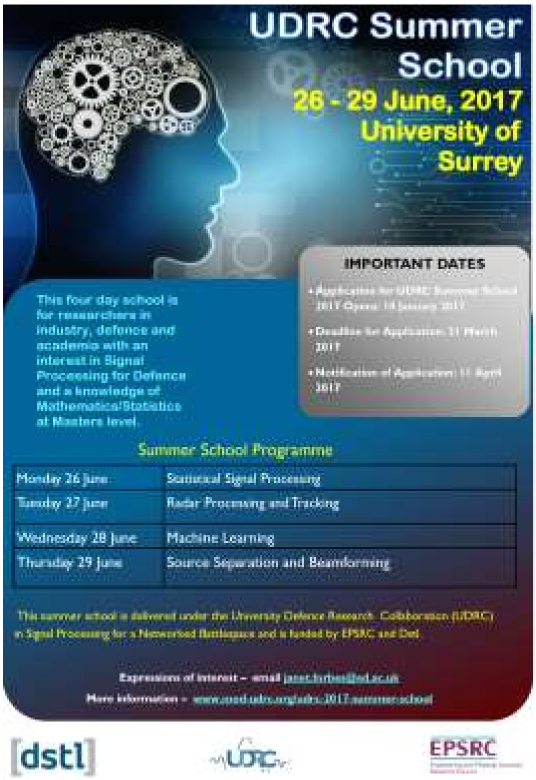

The Summer School in 2017 will take place at the University of Surrey from the 26th to 29th June. Regsitration is open until 31st March and already we have 60 applicants.

The programme has changed slightly in order to provide a radar component for a course that Dstl and industry are running and because of this we expect the applications and attendance from Dstl and industry to increase this year.

Sensor Signal Processing and imaging AIMDay 2016

This AIMday took place on the 21st September and we have had 6 questions posed to our UDRC academics from four companies. We had 35 academics debating the solution and investigating a way forward. Out of the 6 questions, 3 questions have been taken forward after the event (see table 3)

Table 3: Showing questions that are being taken forward.

Company	Question
Roke Manor	Recognising whether a signal is a MIMO transmission or not
Thales	Distributed Acoustic Source Localisation with Low Communication Bandwidth
Hook Marine	From the complex motions of a marine vessel at sea, can the roll motion be isolated?

SSPD

SSPD 2016 was a successful conference with 120 attendees; two keynotes speakers; Dr Philip Perconti from US Army Research Laboratory and Prof Geert Leus from the Delft University of Technology, The Netherlands and two invited speakers on radar and tracking, Dr Antonio De Maio and Dr Clark respectively. We accepted 12 oral presentations and 21 poster presentations and were pleased to have representation from a variety of industry and US Air Force and Army.

SSPD 2017 will take place in Savoy Place, London on the 6th to 7th December. This event is partnering with IET Intelligenet Signal Processing Conference which takes place on the 4th-5th December at the same venue. We are currently investigating keynote and inviteed speakers for this event.

UDRC Phase Two End of Project Showcase Event

The UDRC Phase 2 End of Project show case event will take place at Tidworth Garrison Theatre on the 21st February 2018. Both consortia will attend and will showcase their research in a series of presentations, demonstrations and posters.

Publications

[P1] Mehrdad Yaghoobi, Shaun I. Kelly and Mike E. Davies , Autofocusing for High Resolution 3D Synthetic Aperture Radar: a Phase Recovery Approach, under-revision, 2017.

[P2] Mehrdad Yaghoobi and Mike E. Davies, Raman Spectral Decomposition with Nonlinearity Artefacts and Fluorescent Backgrounds, UDRC Sparsity Based Raman Stage 3 Report, ED TIN 2-12, February 2017.

[P3] Di Wu, Mehrdad Yaghoobi and Mike E. Davies, Sparsity Driven Moving Targets and Background Separation via Multi-Channel SAR, submitted to IEEE Transactions of Geoscience and Remote Sensing, November 2016.

[P4] Mike E. Davies , Sparse signal separation and imaging in Synthetic Aperture Radar, CoSeRa, 19-22 September 2016, (plenary talk p62).

[P5] Mehrdad Yaghoobi, Shaun Kelly and Mike E. Davies, Range Focusing in Volumetric SAR: a Phase Recovery Approach, European SAR (EuSAR) Conference, June 6-9, 2016.

[P6] Di Wu, Mehrdad Yaghoobi and Mike E. Davies , A New Approach to Moving Targets and Background Separation in Multi-Channel SAR, IEEE Radar Conference, 2-6 May, 2016.

[P7] Mehrdad Yaghoobi, Shaun I. Kelly and Mike E. Davies ,Phase Recovery for 3D SAR Range Focusing, IEEE Radar Conference, 2-6 May 2016.

[P8] Rhea J. Clewes, Mehrdad Yaghoobi, Di Wu and Mike Davies , Complex Mixture Analysis Using Hand-Held Raman Chemical Detectors and Novel Spectral Deconvolution Algorithms, SciX, 18-23 September 2016 (talk).

[P9] Mehrdad Yaghoobi, Di Wu, Rhea J. Clewes and Mike E. Davies, Fast Sparse Raman Spectral Unmixing for Chemical Fingerprinting and Quantification, SPIE.Security + Defence, 26-29 September 2016.

[P10] Di Wu, Mehrdad Yaghoobi and Mike E. Davies, Digital Elevation Model Aided SAR-based GMTI Processing in Urban Environments, SSPD 22-23 September 2016.

[P11] K. Kim, M. Uney, B. Mulgrew, Detection of manoeuvring low SNR objects in receiver arrays, Proceedings of the SSPD 2016, Edinburgh UK, September 2016.

[P12] K. Kim, M. Uney, B. Mulgrew, Simultaneous tracking and long time integration for detection in collaborative array radars, IEEE Radar Conference 2017, May 2017, to appear.

[P13] K. Kim, M. Uney, B. Mulgrew, Detection via simultaneous trajectory estimation and long time integration, IEEE Transactions on Aerospace and Electronic Systems, draft for submission.

[P14] S. Herbert, J. Hopgood, B. Mulgrew, Exact MMSE adaptive waveform design for active sensing with applications to MIMO radar, IEEE Transactions on Signal Processing, currently awaiting Dstl clearance for submission.

[P15] M. Uney, B. Mulgrew, D. Clark, Latent parameter estimation in fusion networks using separable likelihoods, IEEE Transactions on Signal and Information Processing Over Networks, special issue on inference and learning over networks, under revision.

[P16] M. Uney, B. Mulgrew, D. Clark, `Distributed localisation of sensors with partially overlapping field-of-views in fusion networks, The 19th International Conference on Information Fusion (Fusion 2016), Heidelberg, Germany, July 2016.

[P17] Yan Pailhas, Yvan Petillot, Bernard Mulgrew, Increasing circular SAS resolution via adapted wave atoms deconvolution, Journal of the Acoustical Society of America, accepted, 2016.

[P18] Yan Pailhas, Jérémie Houssineau, Yvan R. Petillot, Daniel E. Clark, Tracking with MIMO sonar systems: applications to harbour surveillance, IET Radar, Sonar & Navigation, DOI: 10.1049/iet-rsn.2016.0080, 2016.

[P19] Yan Pailhas and Yvan Petillot, Spatially Distributed MIMO Sonar Systems: principles and capabilities, IEEE J. Ocean. Eng, DOI: 10.1109/JOE.2016.2593602, 2016.

[P20] Yan Pailhas, Keynote lecture: The underrated phase, Acoustic and Environmental Variability, Fluctuations and Coherence Conference, A. B. Wood Medal Lecture, Cambridge, UK, December, 2016.

[P21] Yan Pailhas, Yvan Petillot, High resolution systems: the resolution problem, invited paper, 172nd Meeting, Honolulu, Hawaii, November, 2016.

[P22] Yan Pailhas, Yvan Petillot, Neither PAS nor CAS: MIMO, Oceans'16 MTS/IEEE, Shanghai, China, April, 2016.

[P23] Yan Pailhas, Yvan Petillot, Circular synthetic aperture sonar image resolution theory, invited paper, Acoustics'17 Boston, Boston, US, June, 2017.

[P24] Anthony Lyons, Roy Hansen, James Prater, Warren Connors, Yan Pailhas, Internal waves effects on seafloor imagery and bathymetry estimates, Acoustics'17 Boston, Boston, US, June, 2017.

[P25] Salvatore Caporale, Yvan Petillot, A Novel Motion Compensation Approach for SAS, SSPD Sensor Signal Processing for Defence, 2016. web: http://ieeexplore.ieee.org/document/7590585/

[P26] J. Bonnel, S. Caporale, A. Thode, Estimation of mode amplitude using warping and phase compensation, The Journal of the Acoustical Society of America, March 2017.

[P27] S. Caporale, From Frequency Warping to Time Warping: just a jump to the left, in Proc. IMA International Conference on Mathematics in Signal Processing, 2016.

[P28] Puneet S Chhabra, Andrew M Wallace and James R Hopgood, Improved Image Discrimination using Fast Non-linear Orthogonal Dictionary Learning, EUSIPCO 2017 (Submitted).

[P29] Puneet S Chhabra, A. Maccarone, A. McCarthy, A. M. Wallace and G. S. Buller, Underwater Photon Counting LiDAR: Data Analysis for Foliage Penetration and MCM, IEEE Oceans, Aberdeen, 2017 (Submitted).

[P30] Puneet S Chhabra, Y. Pailhas, A. M. Wallace, J. R. Hopgood and Y.R. Petillot, Target Classification in SAS Imagery using Orthogonal Basis Selection, IEEE Oceans, Aberdeen, 2017 (submitted).

[P31] Puneet Chhabra, Yan Pailhas, Andy Wallace, James Hopgood and Yvan Petillot, Dictionary merging for simultaneous approximation and discrimination of sonar signals, IEEE Journal of Oceanic Engineering, 2016 (Under Review).

[P32] Puneet Chhabra, Aurora Maccarone, Aongus McCarthy, Andy Wallace and Gerald Buller, Discriminating underwater lidar target signatures using sparse multi-spectral depth codes, IEEE Sensor Signal Processing for Defence, 2016.

[P33] Puneet Chhabra, Aurora Maccarone, Aongus McCarthy, Andy Wallace and Gerald Buller, Novel spectrally enhanced descriptors for discriminating single photon counting lidar data, IEEE Pattern Analysis and Machine Intelligence (PAMI) 2016 (In preparation).

[P34] Cristian Rusu and John Thompson, Learning Fast Orthonormal Sparsifying Transforms, submitted to SPARS 2017.

[P35] S. Caporale, Y. Petillot, A New Framework for Synthetic Aperture Sonar Micronavigation, to be submitted to the IEEE Journal of Ocean Engineering.

[P36] S. Caporale and Y. Petillot, Time Warping and Interpolation Operators for Piecewise Smooth Maps, to be submitted IEEE Trans. on Signal Processing.

[P37] D'Arca, E.; Robertson, N. M.; Hopgood, J.R, Robust indoor speaker recognition in a network of audio and video sensors, Signal Processing, Vol. 129, 12.2016, p. 137-149. 2016.

[P38] Hu, Guosheng; Yan, Fei; Chan, Chi-Ho; Deng, Weihong; Christmas, William ; Kittler, Josef; Robertson, Neil M., Face recognition using a unified 3D morphable model, Proceedings of the 14th European Conference on Computer Vision. Springer Verlag, 2016.

[P39] Baxter, R. H.; Lopez-Guevara, T.; Robertson, N. M., Spatio-temporal anomaly detection in large streaming datasets of target trajectories, Pattern Analysis and Machine Intelligence, IEEE Transactions on. 2016 (In revision).

[P40] Cristian Rusu, John Thompson, Learning Fast Sparsifying Overcomplete Dictionaries, submitted to EUSIPCO 2017.

[P41] Delande, E. D. and Houssineau, J. and Clark, D. E., Multi-object filtering with stochastic populations, Automatic Control, IEEE Transactions on, submitted (arXiv:1501.04671v2).

[P42] Schlangen, I. and Delande, E. and Houssineau, J. and Clark, D. E., A second-order PHD filter with Panjer point process prior, Signal Processing, IEEE Transactions on, submitted.

[P43] Bryant, D. S. and Delande, E. D. and Gehly, S. and Houssineau, J. and Clark, D. E. and Jones, B. A., The CPHD Filter with Target Spawning, Signal Processing, IEEE Transactions on, 2016.

[P44] Nagappa, S. and Delande, E. D. and Clark, D. E. and Houssineau, J., A Forward-Backward Cardinalized Proability Hypothesis Density Smoother, IEEE Transactions on Aerospace and Electronic Systems, 2017.

[P45] Delande, E., Frueh, C., Franco J, Houssineau, J. and Clark D., A novel multi-object filtering approach for space situational awareness, AIAA Journal of Guidance, Control and Dynamics (JGCD), in revision.

[P46] Schlangen, I. and Delande, E. and Houssineau, J. and Clark, D. E., A PHD Filter with Negative Binomial Clutter, Information Fusion, Proceedings of the 19th International Conference on, 2016 A PHD Filter with Negative Binomial Clutter

[P47] Simpson, C. and Hunter, A. and Vorgul, S. and Delande, E. and Franco, J. and Clark D. and Rodriguez Perez, J., Likelihood modelling of the Space Geodesy Facility laser ranging sensor for Bayesian filtering, Sensor Signal Processing for Defence (SSPD), 2016.

[P48] Delande, E. D., Houssineau, J., Franco, J., Frueh, C., and Clark, D. E., A new multi-target target tracking algorithm for a large number of orbiting objects, AIAA 27th Space Flight Mechanics Meeting, 2017 (to appear).

[P49] Schlangen, I., Bahrti, V., Delande, E. D. and Clark, D. E., Joint Multi-Object And Clutter Rate Estimation With The Single-Cluster PHD Filter, International Symposium on Biomedical Imaging (ISB 2017), 2017 (to appear).

[P50] Franco, J. and Delande, E. D. and Frueh, C. and Houssineau, J. and Clark, D., Probabilistic Orbit Determination Using a Sensor Co-ordinate Parametrization, Journal of Guidance, Control and Dynamics, submitted.

[P51] C. Rusu, N. Gonzalez-Prelcic and R. W. Heath, Fast Orthonormal Sparsifying Transforms Based on Householder Reflectors, IEEE Transactions on Signal Processing, 64 (24), pp. 6589-6599, 2016.

[P52] Saurav Sthapit, James R. Hopgood, Neil M. Robertson, John Thompson, Offloading to neighbouring nodes in smart camera network, EUSIPCO 2016, pp 1823-1827.

[P53] Cristian Rusu, John Thompson, Neil Robertson, Balanced Sensor Management Across Multiple Time Instances via L-1/L-Infinity Norm Minimization, to be presented at ICASSP 2017.

[P54] C. Blair, J. Thompson and N. Robertson, Improving Object Detector Algorithms using Uncertainty and Reliability. Journal paper Image and Vision Computing, submitted October 2016.

[P55] Cristian Rusu, John Thompson, Neil M. Robertson, Sensor management with time, energy and communication constraints, Journal Paper to IEEE Transaction on Signal Processing (submitted Feb 2017).

[P56] Cristian Rusu, John Thompson, Learning Fast Sparsifying Transforms, Journal to IEEE Transactions on Signal processing (under revision for second round review March 2017).

[P57] Cristian Rusu, John Thompson, Power Reduction on Image Processing on FPGAs (submitted to IEEE CSVT) Feb 2017.

[P58] Cristian Rusu, John Thompson, On the use of tight frames for optimal sensor placement in time-difference of arrival localization, submitted to EUSIPCO 2017.

[P59] Paulo Garcia, Deepayan Bhowmik, Robert Stewart, Andrew Wallace and Greg Michaelson, "Optimal Memory Allocation and Power Minimization for FPGA-Based Image Processing", IEEE Circuits and Systems for Video Technology (submitted Feb 2017).

[P60] Stewart, Robert; Michaelson, Greg; Garcia, Paulo; Bhowmik, Deepayan; Wallace, Andrew; "RIPL: A Parallel Image Processing Language for FPGAs", ACM Transactions on Reconfigurable Technology and Systems (submitted Feb 2017)..

[P61] Robert Stewart, Greg Michaelson, Idris S.Ibrahim, Paulo Garcia, Andrew Wallace, Bernard Berthomie, "Enabling Parallel Dataflow Graph Transformations using Petri Net Analysis", 10th International Conference on Graph Transformation(submitted Feb 2017) (ICGT 2017).

[P62] Idris S.Ibrahim, Robert Stewart, Greg Michaelson, Paulo Garcia, Andrew Wallace, "An Interactive Visual Dataflow Tranformation Framework", 20th International Workshop on Software and Compilers for Embedded Systems (submitted Feb 2017).

References

[R1] M. Uney, B. Mulgrew, D. Clark, `Distributed estimation of latent parameters in state space models using separable likelihoods, in the Proc. Of the 41st ICASSP 2016, Shanghai, China, March 2016.

[R2] M. Uney, B. Mulgrew, D. Clark, `A cooperative approach to sensor localisation in distributed fusion networks, IEEE Transactions on Signal Processing, March 2016.

[R3] H. Godrich, A. Haimovich, and R. Blum, Target localization accuracy gain in mimo radar-based systems, IEEE Transactions on Information Theory, vol. 56, no. 6, pp. 2783–2803, June 2010.

[R4] R. Niu, R. Blum, P. Varshney, and A. Drozd, Target localization and tracking in noncoherent multiple-input multiple-output radar systems, IEEE Transactions on Aerospace and Electronic Systems, vol. 48, no. 2, pp. 1466–1489, April 2012.

[R5] W. Huleihel, J. Tabrikian, and R. Shavit, Optimal adaptive waveform design for cognitive MIMO radar, IEEE Transactions on Signal Processing, vol. 61, no. 20, pp. 5075–5089, Oct 2013.

[R6] L. Jiang, S. S. Singh, and S.Yildirim, Bayesian tracking and parameter learning for non-linear multiple target tracking models, IEEE Transactions on Signal Processing, vol. 63, no. 21, pp. 5733–5745, Nov 2015.

[R7] M. Uney, M. Cetin, Monte Carlo optimisation of decentralised estimation networks over directed acyclic graphs under communication constraints , IEEE Transactions on signal processing, vol. 59, issue 11, pg. 5558-5576, November 2011.

[R8] B. Mu, G. Chowdhary, J. P. How, Efficient distributed sensing using adaptive censoring-based inference, Automatica, vol 50, issue 6, pg. 1590-1602, June 2014.

[R9] A. Lepautre, O. Rabaste, F. Le Gland, Multitarget likelihood computation for track-before-detect applications with amplitude fluctuations of type swerling 0,1 and 3, IEEE Aerospace and Electronic Systems, Vol 52., Issue 3., June 2016.

[R10] Houssineau, J.; and Clark, D. E., A sequential Monte Carlo approximation of the HISP filter, European Signal Processing Conference, 2015.

[R11] Morris, B. T.; Trivedi, M. M., Trajectory Learning for Activity Understanding: Adaptive Approach, IEEE Transactions on Pattern analysis and Machine Intelligence, 33(11), pp 1-17, 2011.

[R12] Laxhammar, R.; Falkman, G., Online learning and sequential anomaly detection in trajectories, IEEE Transactions on Pattern analysis and Machine Intelligence, 36(6), pp 1158-1173, 2014.

[R13] Baxter, R. H.; Robertson, N. M., Lane, D. M., Human behaviour recognition in data scarce domains, Pattern Recognition, 48(8), pp 2377-2393, 2015.

[R14] Xiao, T.; Li, S.; Wang, B.; Lin, L.; Wang, X., End-to-End Deep Learning for Person Search, *ArXiv*., 2016.

[R15] Zheng, L.; Huang, Y; Lu, H.; Yang, Y., Pose Invariant Embedding for Deep Person Re-identification, CVPR, 2017.

[R16] Zhang, K.; Zhang, Z.; Li, Z.; Member, S.; Qiao, Y.; Member, S., Joint Face Detection and Alignment using Multi-task Cascaded Convolutional Networks, SPL, (1), 1–5, 2016.

[R17] Bertinetto, L.; Valmadre, J.; Henriques, J. F.; Vedaldi, A.; Torr, P. H. S., Fully-Convolutional Siamese Networks for Object Tracking, *ArXiv*., 2016.

[R18] Delande, E. and Houssineau, J. and Clark, D. E., Performance metric in closed-loop sensor management for stochastic populations, Sensor Signal Processing for Defence (SSPD), 2014

[R19] Houssineau, J., Representation and estimation of stochastic populations, PhD Thesis, Heriot-Watt University, 2016.

[R20] Delande, E., Üney, M., Houssineau, J., and Clark, D. E., Regional variance for multi-object filtering, IEEE Transactions on Signal Processing, 2014.

[R21] Andrecki, M. and Delande, E. D. and Houssineau, J. and Clark, D. E., Sensor management with regional statistics for the PHD filter, Sensor Signal Processing for Defence (SSPD), 2015.

[R22] Mahler, R. P. S., Statistical Multisource-Multitarget Information Fusion, Artech House, 2007.

[R23] Franco, J. and Delande, E. D. and Frueh, C. and Houssineau, J. and Clark, D. E.,
A Spherical Co-ordinate Space Parameterisation for Orbit Estimation, IEEE Aerospace conference , 2016.

[R24] Hagen, O. and Houssineau, J. and Schlangen, I. and Delande, E. D. and Franco, J. and Clark, D. E., Joint Estimation of Telescope Drift and Space Object Tracking, IEEE Aerospace conference, 2016.

[R25] Pak, A. and Clark, D. E. and Correa, J. and Adams, M. and Delande, E. D. and Houssineau, J. and Franco, J. and Frueh, C., Joint Target Detection and Tracking Filter for Chilbolton Advanced Meteorological Radar Data Processing, Advanced Maui Optical and Space Surveillance Technologies Conference, 2016.

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

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

Dr Neil Robertson, Work Package 4 leader N.Robertson@qub.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

Dr Merhdad Yaghoobi, Work Package 1 M.Yaghoobi-Vaighan@ed.ac.uk

Prof Andrew Wallace, Work Package 3, 4 & 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 & 5 M.Sellathurai@hw.ac.uk

Project Management

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