8th IMA Conference on Mathematics in Defence and Security Mathematics for Decision Support

Abstracts

15th October 2025

INSTITUTE OF
MATHEMATICS &
ITS APPLICATIONS

Contents

Which data do we analyse further to inform where to look for MH370?

Prof Simon Maskell

MH370 went missing on 8 March 2014 with 239 people onboard, travelling for several hours after its ADS-B transponder stopped transmitting the aircraft's position. Data from several sources have been used by multiple teams to attempt to localise the aircraft's current position. One might think that the decision of where to search can then be straightforwardly posed in terms of fusing all the data together to define a search area. However, any such fusion process needs to make explicit any assumptions about any decisions that were made by anyone that might have been in control of the aircraft. Furthermore, several of the teams who have tried to help find MH370 differ in terms of what assumptions they have decided to make. This appears to result in their conclusions contradicting one another's. There has also been a desire for credible new evidence as a pre-requisite for the search being continued. This emphasises the need to think more explicitly about the long-term impact of any decisions related to which data to analyse on the chance of finding MH370. This talk will consider the search for MH370 through the lens of the decisions in an uncertain past, present and future and explore how these decisions are playing a role in MH370 being found.

Regularised Volumetric SAR

Dr Joshua Hellier, University of Manchester

In this talk we present joint work done by the University of Manchester in collaboration with Dstl and QinetiQ. Synthetic volumetric SAR measurement data for various scenes and apertures were produced by QinetiQ, simulating situations where conventional SAR image formation techniques such as backprojection perform poorly. Using the Computational Imaging Library (CIL) and freshly-written SAR modelling software, UoM produced reconstructions of the imagery by leveraging three well-known regularisation schemes (Tikhonov, L1 and Total Variation). We will discuss the methods of calculation, computational complexity, and qualitative performance, and compare with the equivalent backprojected imagery.

Application of mutli-target tracking and fusion algorithms in a live sensor network

Mr Luke Storey, University of Edinburgh

Target tracking and multi-sensor data fusion remain critical components in modern defence systems, enabling accurate situational awareness in complex and dynamic environments. This presentation outlines the application of contemporary target tracking and track-to-track fusion algorithms to real-world radar sensor data collected during a recent field experiment. The experimental setup involved two radar platforms operating in a coordinated sensing environment, generating simultaneous data streams, sampling rates, and detection characteristics. Two cooperative seaborne rigid inflatable boats (RIBs) were operating in the scene performing a variety of challenging tracking maneuvers. We applied a Kalman filter based joint probabilistic data association model with message passing for multi target tracking to process the radar detections and estimate target kinematics live, in real time. A plot fusion method was used to fuse the two incoming data streams into a single tracking environment. Following the experiment, track-to-track fusion techniques were implemented to combine independent tracking outputs from multiple sensors. An unknown source of interference was present throughout the experiment, increasing the false alarm rates in each sensor periodically. We investigate how track to track fusion techniques can reduce the amount of false tracks generated from large amounts of clutter.

Practical Sensor Management for Defence

Miss Nikki Perree, Dstl

This presentation focuses on the development of sensor management algorithms for applications in defence relevant scenarios with the aim to mitigate practical implementation challenges.

Sensor management is the process of deciding and executing the actions that a group of sensors will take, based on optimising some reward. The development of enhanced sensor management algorithms is essential for achieving pervasive, full spectrum, multi domain ISR. This is one of the five capability challenges identified by the Ministry of Defence in the Science and Technology Strategy (2020)[1]. The defence landscape is evolving rapidly and the requirement for autonomous sensor tasking in contested environments is only increasing.

The practical challenges of applying sensor management algorithms to real world problems include combinatorics, short-sightedness (myopia), and balancing multiple (often conflicting) objectives. We show this with simple examples and then with increasing complexity. We also consider how to incorporate information about the environment in which the sensor is operating. For example, how objectives may vary based on areas of interest, areas of high threat, or obstacles and occlusions.

Experimentation utilises Stone Soup, an open-source framework for Bayesian state estimation[2]. Primarily, Stone Soup is for building, testing, and evaluating tracking algorithms. Recent additions enable users to integrate sensor management algorithms into their applications. Through a series of case studies and simulations, we show how the components in Stone Soup can be assembled in a way which can be tailored to a sensor management scenario.

Understanding how these algorithms can practically and effectively be applied to the most relevant defence problems is essential for future implementation.

[1] MOD Science and Technology Strategy 2020 - GOV.UK

[2] Welcome to Stone Soup's documentation! — Stone Soup 1.6 documentation

Discrete Event Simulation modelling in support of data-based decision making

Contributing Author(s): Andrew Pearson (CDS – Defence and Security), Ryan Griffin (CDS – Defence and security)

The ultimately the goal for any defence procurement or maintenance project is to provide equipment to the Ministry of Defence (MoD) giving maximum capability / availability for minimum cost. These two key parameters are generally in opposition and therefore reaching the optimum solution is key to sustained success. This solution needs to be effective over both a sustained period and bursts of increased utilisation.

There are many interconnected parts to this process, one key element is how many spares are needed and where they are held. Maintenance of a system, either planned or reactive is dependent on having the correct spare part, at the correct location, at the right time. There are tools which are used to provide recommendations, but these are based on averages and often do not give enough confidence to cover most possible situations.

This presentation shows how a digital twin and discrete event simulation model of the support solution for a fleet of vehicles has been used to provide data driven decisions allowing a force commander to deploy and sustain equipment required for operations. This model is used to improve the availability of a fleet of different variants of a vehicle and therefore enhance the capability of the overall system, whist being cognitive of the incurred costs. By using both predicted information, such as failure rates and actual in-service data, the model can provide the decision makers up-to-date recommendations allowing spares to moved dynamically, thus avoiding any unnecessary drops in availability.

Uncertainty Analysis in Experimental Vibration Datasets

O. Carter, N. Turner

Vibration analysis underpins the capability for Condition based Monitoring in Maritime, Land and Air defence domains. Understanding structural dynamics of in-service platforms is important as it allows for non-invasive fault monitoring of mechanical deterioration over the lifespan of the platform. Vibration datasets for determining mechanical impedances can be gathered in-situ either via manual force impulse excitation or under operational excitation from machinery forces. There is a need to understand how experimental errors propagate when utilising these measurements for forward calculations since errors from inverse modelling may dominate predictions. In this contribution, uncertainty analysis is performed on a simple numerical model, in which force excitations are measured at specified response locations. A range of model parameters are investigated such as measurement sampling density, measurement location and structural damping in the presence of varying background noise levels (specific to measurement channels and at a system level). This supports the understanding of the propagation of noise through inverse and forward calculations for mechanical impedance measurements. A method for mitigating the uncertainty in these measurements is also discussed using a Principal Component Analysis (PCA) based control approach. This present work enables a quantitative assessment of error propagation through measurements and an approach to mitigate these errors, enabling more effective utilisation of vibration datasets in noisy environments.

The Importance of Operational Context

Dr Phillippa Payne, BAE Systems

“If, instead of sending the observations of able seamen to able mathematicians on, land, the land would send able mathematicians to sea, it would signify much more to the improvement of navigation and to the safety of men’s lives and estates on that element.” Sir Isaac Newton.

Mathematics is fundamental to success in security and defence and has been throughout history. Solving operational problems through mathematics has been the duty of many of us over the years and continues to highlight the importance of the multidisciplinary approach of operator and mathematician as a team. This teamwork is essential for bridging the gap between abstract theory and practical application. It can be argued that operational context provides a grounding to mathematical applications and shows how it contributes to social and scientific goals. This talk explores historic and recent examples of the application of mathematical knowledge to operational problems and poses the question how do we ensure the community’s continued relevance and operational knowledge of research, security and defence and adapt to greater flexibility in team constructs.

Uncertainty Quantification of a computationally expensive model used in defence

Dr Ed Ryan (Atomic Weapons Establishment)

In many situations in defence, mathematical models are used to simulate a real-life event. For example, we may wish to use such a model to predict the effectiveness of a system following a military event. As well as getting a point estimate of the model's output, modellers are increasingly interested in quantifying the uncertainty on this output and carrying out other statistical analyses such as global sensitivity analysis. However, this typically involves running the model tens of thousands of times. When the model is computationally expensive, we instead carry out these model realisations using a faster running surrogate model which is a statistical model that accurately maps the inputs to the outputs of the simulator model. Over the past decade there has been tremendous progress in the development in different types of surrogate models, but there is less understanding about which type of surrogate model is most appropriate to different types of simulator models. In this talk, I will present results from three commonly used surrogate models – based on Gaussian processes, deep learning and random forests.

Optimising Defence Coverage with Genetic Algorithms

Semonte Mansur
Alice Sendall

Defence operations frequently demand rapid, data-driven decision-making. A prominent challenge is combinatorial optimisation – for example, effective deployment of limited military resources, such as surveillance sensors, to maximise strategic performance metrics within operational limits. This can be framed as a Maximum Coverage Problem (MCP).

In this presentation, we explore a generalised MCP formulation based on real-life defence and military scenarios. The problem is modelled using overlapping polygons, with each polygon representing an area of influence, such as sensor coverage; the overlaps themselves represent a combined coverage. We seek to optimise a selection of overlaps, and/or groups of overlaps, such that we maximise a defined performance metric (e.g., total coverage) while adhering to budgetary or operational constraints.

We solve this problem using a genetic algorithm, a method inspired by natural selection that iteratively refines candidate solutions, under fitness criteria informed by strategic objectives. For security reasons, we are not able to provide specific details of the context. However, we demonstrate the approach and solutions based on simulated applications that reflect realistic defence and military planning. This work highlights how optimisation using genetic algorithms can enhance operational effectiveness and decision-making in complex scenarios.

An Adaptive Anti-GNSS Spoofing Framework for Autonomous Maritime UAV Surveillance Operations

Jack Carmichael, Yoann Altmann, James Hopgood

GNSS spoofing attacks present a significant threat to both civilian and military platforms by intentionally feeding the operator false information about their position and reducing the chance of mission success. Autonomous systems, which lack human intuition to detect when they're being misled, are particularly vulnerable to this and robust countermeasures therefore need to be considered in their use and design. This work proposes an adaptive multi-platform sensor fusion framework to detect and mitigate GNSS spoofing for autonomous Uncrewed Aerial Vehicles (UAVs) conducting maritime surveillance operations. This framework cross-validates a single resource-constrained ship-based radar tracking information, that has other duties such as missile detection, with tracking information from on-board autonomous UAV sensor measurements and employs self-consistency diagnostics resulting from likelihood estimates from multi-sensor tracking algorithms to detect and mitigate GNSS spoofing attacks. When anomalies are detected by either agent, the system dynamically adapts the ship-to-UAV tracking intervals and the UAV on-board self-checking parameters to validate if an attack is taking place, and attempts to continue operation by relying on derived positional estimates rather than a potentially compromised GNSS. This framework is designed for a surface marine environment where signals of opportunity and geographical features may be sparse, limiting the use of feature recognition systems, and also account for common sensor fusion problems such varying data rates. The approach is modular and lays the groundwork for future experiments with varying sensors and platforms to combat counteract GNSS spoofing for both civilian and military applications.

Outplace: Interactive Games for Complex EME Optimisation

Emma Halliday, Allan Player

Outplace is a framework for abstract adversarial games in the Electromagnetic Environment (EME) developed under DSTL's Integrated Electromagnetic Activities Options (IEMAO) project to explore optimisation algorithms and decision support for Electromagnetic Activities (EMA).

Optimisation for EMA is a multi-agent, high dimensional problem dependant on physical space/time as well as the spectral properties of power, frequency, and polarisation. Providing decision support across an EME will require error propagation methods and uncertainty quantification.

Outplace presents the opportunity to investigate and develop theoretical techniques for optimisation and decision making across an EM Enterprise; a complex EME, with interconnected data sources, sensors, effectors, and platforms.

The purpose of this work is to develop and demonstrate the skills required to rapidly deploy prototype decision support tools, using optimisation and data science to understand and navigate the future EME.

It challenges participants to engage in tactical planning and adversarial thinking, offering insights into the interplay between human decision-making and automated optimisation tools.

An example problem for Outplace is; a Blue team places static EW sensors to maximise their line-of-sight coverage while the Red team chooses UAV search routes to find and neutralise these sensors. This raises interesting questions including:

- 'Where would the Red team be likely to search first?'
- 'How should assets be deployed when different terrain types have different search durations within a time-limited scenario?

Since Outplace is a flexible code-based framework, any type of Electromagnetic Operations (EMO) challenge can be presented for decision-support tools and optimisation algorithms to be developed against.

There is potential to run an interactive demonstration for the audience on a web app to stimulate discussions on optimisation methodologies, adversarial thinking, and automation's role in EMO decision-making.

Abstract for 8th IMA Conference on Mathematics in Defence and Security

Iain High

Automatic Modulation Classification is the task of identifying the modulation scheme used to encode information onto a carrier for an incoming transmission. This is foundational for further signal intelligence tasks in both military and civilian contexts. Current approaches fail to address the challenges of wideband radio frequency environments that simultaneously contain many signals, and previously unseen modulation schemes, which limits their real-world use cases. This research leverages computer vision techniques to address these challenges. More specifically, the You Only Look Once single-shot object-detection framework was employed to perform wideband modulation classification with open-set recognition in a single forward pass. The proposed model processed wideband signals, partitioned the spectrum into fixed sub-bands and predicted, for each sub-band, frequency offsets, bandwidths, confidence scores, and modulation class probabilities. The model was trained and evaluated on a synthetically generated dataset comprising multi-carrier transmissions with up to four signals per frame. The dataset covered common digital modulation schemes under additive white Gaussian noise and multipath fading channels across a signal-to-noise ratio range from -20 dB to 20 dB. Preliminary results showed that the proposed approach achieved classification accuracy comparable to state-of-the-art narrowband closed-set models while providing wideband capability and open-set flexibility. By addressing wideband operation and unknown modulation schemes, the framework can potentially improve downstream tasks such as adaptive jamming countermeasures and signal decryption. Further, it lays the groundwork for more resilient civilian spectrum monitoring, interference management, and dynamic spectrum access, helping to secure and optimise the radio-frequency ecosystem.

RF Localisation in Infrastructure-limited Environments through Machine Learning

James Hollingdale1,2, Prof. John Thompson1, Dr. João Mota2
1Institute for Imaging, Data and Communications (IDCOM), School of Engineering, The University of Edinburgh, Kings Buildings, Edinburgh EH9 3BF, UK.
2Institute of Sensors, Signals & Systems (ISSS), Heriot-Watt University, Edinburgh EH14 4AS, UK.

Position localisation is fundamental for modern Radio Frequency (RF) communication systems and is relied upon for many critical services. Present line-of-sight localisation systems can consistently achieve <1m localisation accuracy. However, in complex scenarios such as in defence and security, operational environments can often lack established infrastructure or reliable line-of-sight (LOS) conditions. Non-line-of-sight (NLOS) scenarios can significantly degrade localisation accuracy due to errors in channel parameter estimation, thereby posing significant challenges for conventional localisation methods. Emerging Machine Learning (ML) techniques offer improved performance in complex NLOS environments, while in addition acting as a surrogate model that can improve computational efficiency. However, there has been little research into how ML-based localisation methods can be transferred effectively to dedicated hardware platforms to leverage this increased efficiency and enable new applications. This is particularly crucial for low size, weight and power (Low-SWAP) systems that are increasingly important for defence scenarios.

This PhD research project seeks to address these limitations through the development of MLbased localisation methods optimised for deployment on dedicated hardware, such as FieldProgrammable Gate Array (FPGA), Graphics Processing Unit (GPU) and System-on-a-Chip (SoC) platforms. This work seeks to build on two major improvements observed in the existing ML literature. First, that model-driven ML methods enable the creation of lightweight architectures, which in-turn provide improvements in computational efficiency[1]. Second, that ML-enhanced localisation pipelines, such as those involving enhanced information sharing [2], can provide a greater robustness and flexibility to localisation networks. This talk will discuss how advances in either of these two areas will allow for more robust localisation systems that will be crucial to the defence sector in future. In particular, they will support the rapid deployment of temporary local networks in unfamiliar or hostile environments that can dynamically aid in operational decision-making.

Acknowledgment:
This work was supported by the Engineering and Physical Sciences Research Council and Ministry of Defence

Centre for Doctoral Training in Sensing, Processing and AI for Defence and Security, [EP/Y013859/1].

References:

[1] S. Brennsteiner, T. Arslan, J. S. Thompson and A. McCormick, "LAMANet: A Real-Time, Machine LearningEnhanced Approximate Message Passing Detector for Massive MIMO," in IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 31, no. 3, pp. 382-395, March 2023, doi: 10.1109/TVLSI.2022.3225505.

[2] F. K. Coutts, J. Thompson and B. Mulgrew, "Flexible Compression for Efficient Information Sharing in a Network of Radio Frequency Sensors," in IEEE Transactions on Radar Systems, vol. 3, pp. 332-348, 2025, doi:
10.1109/TRS.2025.3529760

Sensor Management in Stone Soup

Nicola Perree

Sensor management is the process of deciding what actions sensors will take, based on optimising against Intelligence, Surveillance and Reconnaissance (ISR) requirements. It is a well-established area of research in academic literature. Autonomous sensing systems are widely available commercially for both civil and defence applications. For practical experimentation and testing against use case scenarios there is a need for a reusable, extensible repository of sensor management algorithms.

This poster showcases the sensor management algorithms implemented in Stone Soup[3]. Stone Soup is an open-source Python framework for Bayesian state estimation. The inclusion of sensor management in Stone Soup has enhanced the capabilities of the framework with a suite of sensor management algorithms for use by the sensing and fusion communities. By combining sensor management with tracking algorithms, this enables exploitation and optimisation of sensor taskings to improve situational awareness.

The poster explains how sensor management algorithms are structured in Stone Soup, to fit the modular, object-oriented structure. We highlight some of the examples and applications that have been considered so far. For example, Bayesian search, source term estimation and maritime situational awareness. The methods implemented include greedy algorithms, optimisation approaches and implementations of reinforcement learning. Most reward functions are based on information theory, and the poster shows how users can develop their own reward functions for a specific use case or to achieve certain behaviours.

This work supports the development of real-world sensor management solutions and contributes to future defence research and experimentation.

[3] Hiscocks, S., Harrald, O., Barr, J., Perree, N., Vladimirov, L., Harris, M., Webb, G., Glover, T., Green, R., Rosoman, O., Rogers, E., Dorrington, I., Wright, J., Pike, S., Acuto, A., Hunter, E., Fraser, B., Pritchett, H., Osborne, J., … Wheelhouse, E. (2025). Stone Soup (v1.6). Zenodo. https://doi.org/10.5281/zenodo.15013693

Bayesian Learning of Trajectory Behaviour with a Stochastic Generative Model

Christian James Pollitt | cjpphd@liverpool.ac.uk |

This work is centred on Bayesian computations for learning the parameters that govern the behaviour of a stochastic trajectory model. This generative model is based on Ornstein-Uhlenbeck processes and adopts their mean-reverting behaviour into kinematics. Specifically, we propose a Gibbs-within-Markov Chain Monte Carlo sampler that explores the likely configurations of six parameters specifying the behaviour of a trajectory segment. These parameters specify a mean velocity vector, reversion rates and process noise variances, and our goal is to learn their distribution for a set of trajectories by sampling from the parameter posterior. We use a combination of Gaussian and truncated Gaussian priors and proposal distributions to handle both unconstrained and non-negative parameters and guide inference toward plausible movement dynamics. This sampler can be used to infer parameters of real trajectories for different vessel types recorded using, for example, Automatic Identification System (AIS) messages. We investigate the efficacy of the resulting scheme in learning behavioural models of trajectories empirically using simulations and the mean squared error of the estimated parameters. The resulting Bayesian computation scheme is relevant to the defence and security domain in finding data-driven classification posteriors for trajectories of unknown objects, which also provides uncertainty quantification. These characteristics are designed to lead to improved ability to deal with deceptions and complex maneuvers. The proposed method is sensor-agnostic and can be used with radar, EO/IR, or other sensor data attributing potential power to discern between low-flying aerial trajectories and maritime trajectories that may otherwise be difficult to classify due to low-resolution, surface reflections, and high background noise.