Edge computing for earth observation

Abstracts

SAR satellite clusters from design to exploitation

Victoria Nockles, The Alan Turing Institute

With increasing availability of COTS solutions for cube satellites and recent advances in UAV technology, research into understanding how multiple platforms can work together in a cluster is growing. This brings unique challenges for competing requirements in platform design, tasking, collection, processing, exploitation and dissemination (TCPED). It also brings opportunities for new avenues of research in digitisation, waveform design, synchronisation, communication, configuration design, and processing and analysis. This talk gives an overview of how different applications of Synthetic Aperture Radar are driving platform design for earth observation. I discuss why limitations in power and bandwidth necessitate collaboration with specialists in distributed systems, and touch on the use of novel computer architectures for onboard processing and machine learning.

Machine learning analysis of geodetic data for earthquake and landslide detection

Richard Walters, University of Durham

We are currently in a golden age of tectonic geodesy, ushered in by the recent step-change in volumes and frequency of acquisition of global geodetic data. These data include GNSS and satellite radar observations of active tectonic regions that enable measurements of small ground movements and modifications of the ground surface at high spatial and temporal resolution. Here I'll show current efforts to harness the power of machine-learning approaches in order to fully exploit these rich datasets for a number of applications. These include the isolation of elusive and exotic 'silent earthquakes' in GNSS time-series and 3D time-stacks of radar images, and the rapid detection of earthquake-triggered landslides following a major seismic event from the integration of outputs from multiple radar-based predictive models.

Application of machine learning to detect ground deformation from volcanoes and urban sources in InSAR data

Pui Anantrasirichai, University of Bristol

Satellite interferometric synthetic aperture radar (InSAR) can be used for measuring surface deformation for a variety of applications. Recent satellite missions, such as Sentinel-1, produce a large amount of data, meaning that visual inspection is impractical. Here we use deep learning, which has proved successful at object detection, to overcome this problem. Initially we present the use of convolutional neural networks (CNNs) for detecting rapid deformation events, which we test on a global dataset of over 30,000 wrapped interferograms at 900 volcanoes. We compare two potential training datasets: data augmentation applied to archive examples and synthetic models. Both are able to detect true positive results, but the data augmentation approach has a false positive rate of 0.20% and the synthetic approach has a false positive rate of 0.04%. Then, I will present an enhanced technique for measuring slow, sustained deformation over a range of scales from volcanic unrest to urban sources of deformation such as coalfields. By rewrapping cumulative time series, the detection performance is improved when the deformation rate is slow, as more fringes are generated without altering the signal to noise ratio. We adapt the method to use persistent scatterer InSAR data, which is sparse in nature, by using spatial interpolation methods such as modified matrix completion Finally, future perspectives for machine learning applications on InSAR data will be discussed.

Big data processing of sentinel-1 InSAR to monitor earthquakes, volcanoes and building subsidence

Tim Wright, University of Leeds

The European Commission's two-satellite Sentinel-1 constellation has global coverage with a repeat cycle of 12 days, and European coverage with a repeat cycle of 6 days. Using the JASMIN super-computer, co-located with the UK's archive of Sentinel data at CEDA, we are processing all Sentinel data, which includes 55,000 acquisitions in our areas of interest from 2015 to present, covering 1430 frames, each 250 x 250 km in size. The total unprocessed data volume in our large AOI is ~550 TB and increases by ~1.6 TB per day. We use this to extract measurements of surface motion with an accuracy of approximately 1 mm/yr. I will show how COMET-LiCS process this data set to cover vast global tectonic and volcanic regions at moderate resolution, and Satsense Ltd process the same data at higher resolution to produce a continuously updated ground monitoring product for the UK.

Multistatic ground-based SAR laboratory investigations

Dan Andre, Cranfield University

Multistatic radar geometries are under consideration by research organisations to provide both pervasive sensing and novel physical characterisation of the environment. These new modalities will provide both technical challenges and opportunities, requiring further investigation. This presentation reports ongoing investigations conducted at Cranfield University's Ground-Based SAR laboratory (GBSAR Lab) into both multistatic SAR interferometry and multi-static through-wall (TW) 3D-SAR imaging. An update will also be provided on the analysis of the DSTL-Airbus LF-SAR datadome airborne collection Bright Sapphire II.

Embedded automatic object recognition for synthetic aperture radar systems

Michael Woollard, UCL

Modern synthetic aperture radar systems are capable of capturing wide swaths at very high resolutions, resulting in an immense amount of raw data to process. Even after image reconstruction, transmitting the complex imagery from the platform to a secondary processing site for object extraction can present significant challenges, especially when the platform is subject to emission control measures. This presentation will address methods for performing automatic object recognition on low SWaP-C systems and the associated challenges with respect to accuracy and robustness.

The Alan Turing Institute

The digitisation of future spaceborne SAR systems
Sam Doody, Airbus

Existing spaceborne Synthetic Aperture Radar sensors consist of a single digital channel and implement beam forming in analogue. Future systems will be driven towards a multiple digital channel approach to break the conventional swath-resolution performance limit and offer simultaneous high resolution and width swath imaging along with other novel modes of operation. This can be seen by the use of multiple digital channels for the next generation European Space Agency Copernicus sensors – Sentinel-1 NG and ROSE-L will both implement tens of digital channels. Digital beam forming or processing on-board is a necessity to reduce the data volume to a manageable volume for downlinking.

Future systems may take the digital beam forming approach a step further through the distribution of the digital channels from a single satellite to multiple satellites. This is achieved through implementation of a cluster of SAR satellites where the data of each satellite can be coherently combined with the data of other satellites such that the cluster can act as a single coherent SAR sensor. Whilst this provides even further performance benefits through enhanced phase centre flexibility, achieving a number of unique modes of operation such as single pass interferometry, it also poses additional challenges since the SAR data is now shared over multiple satellites.

Do switches dream of machine learning? Using in-network computing to accelerate real-time processing
Noa Zilberman, University of Oxford

Earth observation generates vast amounts of data, that needs to be managed, processed and analyzed. However, the rate at which data is generated, and the need to provide near real time analytics, is not sustainable using traditional computing and acceleration technologies.

An emerging research area, In-network computing, tries to attend to these limitations. In-network computing is the offloading of standard applications to network devices, where processing is done as the data traverses the network. This talk will provide an introduction to in-network computing, and present some of the applications already offloaded to the network. The use of in-network computing for inference is one area of growing interest, and the talk will discuss challenges, current solutions, and opportunities.
In-network computing offers immense performance and power saving benefits, and this talk will conclude with a vision of using in-network computing in earth observation.

Optimized resource allocation and computation for distributed sensors
Kin Leung, Imperial

Optimization techniques are often used to allocate and share limited communication and computational resources among competing demands. This talk considers the in-network data processing on wireless sensor nodes where data are collected at multiple distributed sources and aggregated (fused) along the way as they are transferred toward the destination. Finding the optimal degree of data aggregation on the sensor nodes in order to minimize the total energy consumption subject to the data quality constraint is NP-hard. However, for specific parameter settings, we propose a distributed solution framework to achieve the optimal trade-off between communication and computation costs. Future work on use of the proposed framework to distributed data/signal processing problems will also be discussed.

The Alan Turing Institute

A hardware platform for multi modal sensing

Phillip Stanley-Marbell, University of Cambridge

All physical measurements are uncertain and measurements from sensors such as accelerometers, gyros, magnetometers, gas sensors, and so on, are no different. Warp is an open-source multi-sensor hardware platform with multiple redundant sensors for each sensing modality and hardware support for trading sensor precision, accuracy, and reliability for performance and power. The multiple duplicate sensors (e.g., three different 3-axis accelerometers) enable new paths of scientific inquiry into quantifying the uncertainty of measurements and computations on those measurements. In this talk, I will provide a brief overview of the three generations of Warp we have engineered and built over the last three years and will highlight new research results enabled by Warp's unique hardware capabilities.

Emerging connectivity approaches for FPGA accelerators

Suhaib Fahmy, University of Warwick

Field Programmable Gate Arrays (FPGAs) have a long history of use for implementing signal processing operations due to their ability to exploit inherent algorithmic parallelism to deliver high performance with a small footprint. In recent years, architectures have evolved significantly, allowing them to be integrated in a variety of novel ways. Integrated with systems-on-chip, they can serve as tightly coupled accelerators to processors running familiar operating systems and software. Their partial reconfiguration capability can allow for dynamic hardware adaptation at runtime to deal with evolving conditions. And more recently, they have demonstrated strengths in tightly coupling computation with network flows of data. This brief talk will take a tour through these approaches for incorporating accelerators in systems and discuss the benefits and trade-offs involved.

Developing SAR and Machine Learning on Intel FPGAs

Jahanzeb Ahmad, Suleyman Demirsoy, Yaprak Eminaga Intel

This talk will be split into three sections. First, we will cover how the latest FPGA devices are architected to help with the ever-growing compute complexity of vision and AI algorithms. In the second part of the talk, we will touch on the high-level development methodologies that help to efficiently code and verify your FPGA designs. In the final part of the talk, we will discuss embedded vision reference designs with particular emphasis on SAR and deep learning.

Strategies for efficient implementation of neural networks in hardware

Rob Mullins, University of Cambridge

The compute and memory requirements of machine-learning workloads are typically very high and performance is typically limited by power and memory bandwidth. Fortunately, there are many opportunities for optimising both networks and hardware implementations. This is especially true when given the flexibility of creating custom hardware or using reconfigurable logic (FPGA). This talk will describe our work on automatically generating efficient Convolutional Neural Networks (CNN) accelerators for FPGAs, promising next steps and broader challenges in this area.