1) Please tell me a little bit about your field of research

Synthetic Aperture Radar (SAR) is one of the most widely-used remote sensing methods in the world, with applications ranging from topography and environmental monitoring to battlefield surveillance. SAR phenomenology often makes analysis of the resultant imagery a non-trivial task, particularly when the collection is performed under poor operating conditions. Effects such as layover, shadowing, and artefacts from multipath can result in an image which bears little resemblance to an optical counterpart. Automatic Target Recognition (ATR) approaches are used to facilitate analysis by highlighting objects of interest for further investigation and/or permitting automatic filtering of candidate targets. For the last few years, I have been developing ATR methods for X-band surveillance systems, with a particular focus on vehicular targets.

My current research project focusses on exploiting hyperfine SAR imagery, such as that obtained from the Bright Spark system (a Thales UK – Dstl joint technical demonstrator), with resolutions of the order of a few centimetres. At these resolutions, an incredible level of detail about a target can be extracted. However, the artefacts previously mentioned usually obscure a large portion of this information from a visual inspection of the image, necessitating the use of ATR and advanced signal processing to interpret it.

2) Can you describe the background to the work that is presented in your Electronics Letters submission?

The simulation pipeline presented in the Letter was originally developed to support my undergraduate research project, which investigated the implementation of ATR in stealth-capable multistatic SAR systems for battlefield surveillance. In particular, the case of small UAV platforms with a tightly-constrained power budget, and correspondingly low onboard computational power, was considered. The classification results shown in the Letter followed from a series of experiments on the simulated data to identify suitable lightweight ATR approaches for rapidly identifying vehicular targets of interest in scenes. Performing these initial detection and classification steps on-board allows more efficient use of communications bandwidth, and a significant reduction in the electromagnetic footprint of the system. In a cognitive radar system, situational information gleaned from this analysis could also be used to modify the surveillance node's behaviour, thus enabling a greater degree of autonomous operation.

3) What is the main advance you have reported in your Letter and what is the significance of this advance?

The classifier presented is one of the fastest convolutional neural network (CNN) approaches to the problem of object detection in detected SAR imagery, capable of processing over 150 images per second with GPU support. It is also worth noting that the high throughput of the approach demonstrated in this Letter enables fast development of such approaches. This methodology is invaluable for rapidly evaluating hypotheses about the behaviour of classifiers during proof-of-concept development.

Where activation analysis and related techniques would be both time consuming and computationally intensive, tailored test targets can be modelled using an appropriate CAD program and a large number of images simulated in a short period. Observing the behaviour of the classifier against these inputs can give important insights into features which are under-represented in the training data. Furthermore, an existing classifier can be quickly tested against novel or confuser targets to validate its in-field performance.

Very little SAR imagery of vehicular targets exists in the public domain, particularly where current-generation armour and other sensitive targets are concerned. The US Air Force Research Labs MSTAR database is the notable exception to this rule, and has informed research in the ATR domain for over two decades. However, it is a very small dataset by the standards of modern machine learning, and the resolution of the images is not representative of that which can be achieved by many modern SAR systems. The methods presented in the Letter permit the generation of comprehensive SAR imagery datasets for arbitrary targets, and the subsequent development of suitable ATR approaches which can be fine-tuned on sparse real data at a later stage.

4) What challenges did you have to overcome during the research for your Letter?

As previously mentioned, publicly available SAR imagery of vehicular targets is hard to find. Verifying the behaviour of the simulator required many hours building and modifying detailed CAD models for vehicles that appeared in the images we had available. It became necessary to develop many tools to accelerate this process and the analysis of the resultant imagery as the datasets grew to millions of images, particularly whilst fine-tuning the simulation parameters.

The work presented here was a part of a much larger project, which involved the design and construction of an X-band radar system, as well as advanced signal processing modules to support the ATR processing in realtime. With tight deadlines in play, finding enough hours in the day to tinker with non-maximal suppression thresholds and other minutiae of the classifier was only made possible by the dedicated support of our technicians and research colleagues.

5) How much has your research field changed since you began working in it, and how do you think it will develop over the next 10 years?

I believe that the major development in imaging radar over the last few years has been the application of compressive sensing techniques to scene reconstruction; traditionally, radar sensing has been founded on well-established sampling theorems from many decades ago. There have also been exciting developments in high-frequency radar, particularly around the 300 GHz range. I attended a UK symposium on radar sensing (EMSIG) this year at which one academic observed dryly that “One man’s carrier frequency is another man’s chirp bandwidth!”.

With increasing bandwidths and finer resolution, the volume of data that must be processed in a SAR system is always growing. Latency and bandwidth considerations already require that significant pre-processing must be handled at the sensor, a model commonly referred to as edge computing. I think that this will lead to greater focus on the architectures of SAR transceivers and the means by which ATR algorithms can be parallelised. Progress in this area will ensure that future SAR systems meet the ever-tighter SWaP-C constraints, particularly in the case of airborne surveillance.