NuSec Science Network Technical Workshop 18th September 2018 Algorithms for Autonomous Decision Making in Nuclear Security

As part of the annual NuSec (Nuclear Security Science Network) technical workshop on "Algorithms for Autonomous Decision Making in Nuclear Security" which was held on the 18 September 2017 at the Wellcome Trust Centre, London, two break-out sessions took place. The first of these sessions was facilitated by Prof Simon Maskell of Liverpool University and focused on the role of algorithms in future nuclear security. The second of these sessions was facilitated by Dr Matthew Stapleton of AWE. Its focus was finding ways to obtain data to help further develop algorithm research. This document is a brief summary of the key observations and outcomes perceived by those present.

Whilst the discussion during the break-out sessions were diverse, some common themes across both parallel sessions emerged. This paper has been broken down by session for clarity. It concludes with a summary of the overall workshop, with common themes and outcomes.

Session 1.

Research to Data

The implicit assumption of any algorithmic development is that some research has led to some technology being embedded in a device (e.g., for detecting radiation) and that the data generated by that device is accessible in some form or other (i.e., such that it could be processed by algorithms).

Data to Algorithm

The algorithmic landscape is large. It is also described using a different language than that typically used by sensor developers. This complicates interdisciplinary collaboration and motivates the need for training such as; accessible tutorials that explain when Deep Learning is applicable, how to embed Physics into the Deep Learning stack and when a (simpler) shallow approach will be just as useful, archives of algorithms that have historically proven useful and data repositories, for example; the content and context that enables an algorithmic developer to start to think about a problem of interest to the sensor developer. Of course, there is also a need for members of these two communities to have both motivation to engage and to effectively collaborate.

While the science of experimental design is well understood, it is not unknown for a sensor to be developed to measure a phenomenon, but not to be designed with future development toward operational use in mind: just because something can be measured doesn't make it necessarily informative to do so. Similarly, the mathematics of Bayesian statistics necessary to develop algorithms to process data (including unstructured data, e.g., describing the context) with guaranteed convergence properties are well understood. Such algorithms can readily handle instrumental uncertainty (though false alarms are, essentially, inevitable). However, perhaps as a result of some of the issues described above, the appropriate application of these algorithms is less pervasive than one might wish to be the case. Furthermore, algorithms that, for example, consider the potential for future deception when being trained in a benign environment, are not yet mature: it isn't fully understood whether a closed training set (for example, as used by a commercially developed algorithm) would be inherently more or less secure than something developed for a nuclear security application specifically.

There is a tension associated with secondary use (i.e., reusability) of algorithms: naïve application of a convenient implementation rapidly turns from insufficient specialisation into lazy design; however, established implementations are likely to be more trustworthy, future-proof and require lower development costs (and shorter development times).

Algorithm to Product

The existence of an algorithm does not necessarily make it innovative or useful to use such an algorithm or cost-effective to integrate software embodying that algorithm into a broader system. Indeed, as an algorithm migrates towards a product, there is a need to be more explicit about its usage, the decisions that will be influenced by its operational usage and the likely outcomes of these decisions.

Turning a prototype algorithm into a product often requires a computationally efficient implementation and typically demands software that is of a higher quality (maintaining software necessitates higher quality code than simply using the code to support research). There is also a tension between IP being perceived as a necessary component of developing a commercial offering and being a potential national security concern.

Product to Usage

Users of products (i.e., people) need to be confident in the operational utility of an algorithmic system in the real operational world (i.e., not just in the "lab") as well as the provenance of the choice of algorithm and/or any training data used to optimise the algorithm's parameters. This demands that appropriate legal issues are addressed but also that the system undergoes verification and validation. While extensive empirical evidence exists, formal validation is currently challenging for "black box" methods such as Deep Learning. This makes it potentially attractive to intentionally mimic, not aim to outperform, a human analyst. In a nuclear security context, the verification and validation needs to include comparison against meaningful baselines and consideration of black swans. Such black swans can occur both as a result of chance and as a result of an adversary's overt attempts at deception. This is just one reason that it is important to have self-diagnostics such as a "police algorithm" in the system.

While interaction between the algorithmic product and human decision will be improved by appropriate design of the human-machine interface (HMI), there is a wider (under-explored) need to design the composite system to play to the strengths of the both the human and the machine.

Conclusions

Discussion at the algorithmically-focused break-out session has been summarised. It appears to the author that there is a need for sensor and algorithmic developers to work more closely together. It also appears that the discussion highlighted some issues (for example; the need to design algorithms to anticipate potential future deception and the need to explicitly optimise the composite human-machine system) that demand extensions to the state-of-the-art.

Session 2.

Lack of data and its impact on algorithm research:

Whilst it is acknowledged that there is a lot of data across government and industry there is limited awareness of what data is available and to what extent the data is accessible, limitations include commercial / personal sensitivities and general awareness of what exists. As an alternative, generating synthetic data can often be valuable, based on knowledge gained through research into "every" scenario specific to a given model. This would however assume the developer has an awareness of

every eventuality of each specific scenario. This approach is costly and time consuming and not practical in the context of most models. It does also raise the question of how representative this data is of real life and how you validate this.

Establishing data requirements:
The vast breadth of algorithms being used require differing quantities of information and data types to be effective. To sufficiently understand the variations and errors within data a good representative dataset is more informative than a large poor quality inconsistent data set. A fundamental requirement of any dataset is the need to know how it has been constructed, for example is it made up of sub populations or specific to the scenario being solved. It is often the case that companies deliberately alter data to conceal commercial or personal sensitive information prior to releasing it. This can impose bias and lead to drawing incorrect conclusions from the data, creating greater challenges during analysis. In the first instance, understanding the ground truth is important to drawing any conclusion. This does not necessarily mean that absolute knowledge of the context of the data is required, and often by excluding this will help overcome issues of data sharing. This does however drive a deeper discussion about the value of the users' knowledge of the data. Does knowledge of the background motivation for data collection provide more useful information than understanding the context in which it was collated? Which of these provide sufficient information of the risks associated with the analysis being performed?

A recent parliamentary and scientific committee meeting discussed standards for data. It covered how data should be governed and monitored and how the public are dealing with the questions of ethical and social issues raised in the use of personal data. Whilst there are many standards for data and Meta data, they invariably lack clear standards on background, context and legal terms and conditions with regards to use, sharing and collection.

Requirements for data capture need to reflect further algorithmic use. It should contain an understanding of the type of data to test and the type of algorithm suitable. Important questions include; is simulated data suitable, does an algorithm require training data, what is a suitable algorithm to use, for example; Support Vector Machine Learning Algorithms using regression techniques to classify data is often effective for identifying outliers. Requirements capture done badly will hinder effective used of data, done well it will enhance sharing and repeated used of data sets.

Availability of data:
There is a wealth of data in the public domain. Familiar sources include:

- Criminal record / database (limited access / used by government and academia)
- Government trade statistics
- Radiological response IAEA database (high quality data)
- Home office database (unclear how accessible this is).
- FSA open data source.
- Other government owned data such as test data and calibration data (DSTL and NPL).

There is a significant amount of useful data is in existence, owned by companies and embedded in commercial activities (i.e. test data, calibration data). However, to realise this potential a suitable mechanism would need to be established to avoid any breaches of security etc. As an example, Goldman Sachs have created "data lakes" which allows fluid and effective handling of data and provides access to teams across the company from Risk to Technology and Compliance.

Greater awareness of the existing availability from companies would go some way to improve issues surrounding the availability of data and also identify where improvement can feasibly be made to the way it is recorded to improve its quality and benefit further. It is often the case in experimental data that the outputs recorded are adequate for the experiment at the time however greater future analysis opportunities for the data have been missed due to the way results have been recorded.

How to improve data exchange:
In the workshop users of data for algorithm development acknowledged that improvement can be made to the way data is exchanged however no specific recommendations were made. The consensus being that industry and government need to get "better at it". It is recognised that by defining clearly the outcome and not just the problem statement will give a better insight into what is suitable data for a given challenge.

A few key points raised during the final discussion highlighted the benefits of data challenges as an effective way of encouraging and developing innovation, creating a positive environment for networking. Also, no matter how effective an algorithm is, the final decision should always be made by a human and not a machine.

Conclusions
The outcome from the discussion of finding ways to obtain data to drive algorithm development session has been summarised. It appears to the author that there are data sources in the public domain and also valuable data not available to the public. Industry and government need to work better at sharing this. It also appears that the discussion highlighted some issues, for example; the need to structure both the problem and the type of solution being sought from the data.

Workshop summary
The workshop consisted of a diverse range of participants with differing levels of algorithm experience. It highlighted some key issues and common themes in both sessions;

- Developers of algorithms and generators of data are often in different groups. They often have different requirements. By better integrating these groups would mean improved communication and produce long term benefits to both the developers and users of algorithms.
- Requirements capture for developers of algorithms tend to be local and project specific however this limits application of the data. Techniques to capture or allow for future requirements would expand the scope of the work being done and potential work for the future.
- Comprehensively capturing the properties of algorithms and background and context of data will ensure that it is used correctly in its direct project and suitable for future use.
- Opportunities exist in the form of existing data but work needs to be done to make it accessible and available.

The workshop identified that holding data challenges would be a positive step to advance integrated data-algorithm working. This would be a good next step for the NuSec Science network and / or related stakeholders.