Challenge Problems

Artificial Intelligence for Decision Making Initiative
Round Two

1. Multi-Modal Narrative Extraction

Given a multi-modal collection of data (eg. images, text, audio) how can we extract a narrative "through-line" from the data? Technologies such as topic modelling and/or knowledge graph embeddings are increasingly being used to understand large collections of text documents. We seek to extend these ideas into other mediums – namely images, audio and video. This research will investigate how information can be extracted from an increasingly heterogeneous collection of data, and represented in a way that allows a human to seamlessly explore the breadth of data available to them.

2. Reasoning with Uncertainty in Temporal Knowledge

We are looking for ways of reasoning across uncertain, temporal knowledge graphs. In particular, online methods which allow a person to validate hypotheses and information by dynamically adding to the graph, and then exploring any higher order effects introduced into the graph structure. Efficient mechanisms for representing, propagating and resolving uncertainty through multiple steps of inference will be a necessity, as will the ability to quickly include and exclude these hypotheses.

3. Effective updating of deep learning models with limited new data

Deep learning models often have a high compute cost to train, and as such it is not always practical to retrain or fine-tune a model regularly to incorporate updates to the training data set. We are interested in the development of effective strategies for incorporating limited quantities of new data, while not forgetting what has been learnt from the original training dataset. Key considerations here are whether a model can be effectively fine-tuned without the original training set, and without forgetting, or what blending of new/old data is needed to best balance retention of model knowledge with incorporation of new data. A key metric for success in this would be to show an ability to utilise the new data while maintaining equivalent performance against a benchmark.

4. Understanding Linguistic Feature Encoding in Transformers

Considerable attention has been given to recent advances in Natural Language Processing (NLP) tasks through the use of transformer architectures. However, there are still many questions left about exactly what these techniques are modelling. Having a deeper understanding of the way in which these networks model language will provide us with an opportunity to improve these techniques and make the modelling processes more efficient. To this end we are interested in research that is aimed towards understanding how individual linguistic features are treated by these networks – for example, what type of syntactic information is generalised in these networks? How effectively do these models capture referential features? Or what type of phonological information is captured? It is expected that the work will be completed in two or more languages.

5. Word search over Wav2vec 2.0 neural representation

Recent advancements in Transformer models have introduced multiple architectures and platforms towards automatic speech recognition (ASR) such as Wav2vec 2.0, SpeechBrain etc. Moreover, there have been recent open source attempts to train such models that can transcribe low resource languages using openly available datasets. While searching for terms can be performed post-transcription, we seek to explore if these searches can be performed pre-transcription e.g. using neural representations over layers of Wav2vec 2.0 Transformer architecture. It may also worthwhile comparing the complexity between searching terms after using state-of-the-art ASR against searching from representations of such Transformer architecture. Furthermore, we wish to understand how searching over neural representations may differ over two unrelated languages.

6. AI Enabled Portfolio Design for Defence Capability Investment Projects

Defence maintains and develops its warfighting capabilities through ongoing investment in major capital investment projects. Designing effective future Defence capability therefore requires decision-support tools for investment portfolio design. A traditional approach to project portfolio design is typically based on the knapsack formalism applied to optimise the portfolio, which is to maximise the overall portfolio value, subject to the budgetary constraints.

However, one of the most significant challenges to the practical implementation of investment portfolio optimisation is in defining and calculating appropriate value metrics and measures for the proposed investment projects. This is particularly challenging in the Defence context where capability systems are very complex and interdependent, multiples alternative courses of action can be chosen to achieve a desired strategic effect, and the entire Defence capability portfolio must be optimised against an uncertain and evolving strategic and technological environment over long timeframes. Therefore assessment of proposed Defence investment projects often involves human judgment and condensing values judgements to a simple value measure becomes problematic.

This challenge is to suggest concepts for AI-enabled decision support tools to assist with the design of an investment portfolio of effective future Defence capabilities. We are looking for innovative ideas to develop algorithms and approaches to form a data-driven value representation for a project in a Defence capability investment portfolio. Values reflect in the utility of the project in a future strategic Defence context using analysis of publicly available and classified data in a variety of formats. For example, projects could be valued directly on their utility against specific contexts and threats, and the methods of AI could be applied to develop likely adversarial procurement strategies (LAPS). The portfolio would then be optimised directly against these LAPS to maximise Defence's ability to achieve the strategic Defence interests articulated in the Defence Strategic Update 2020.

7. Automatic parsing & annotation of semi-structured documents

A lot of information is present within human-generated textual documents, such as pdf documents. However, having this information in a more structured format is useful for machine analysis and better visualisation of information. This challenge focusses on developing approaches to automatically extract structured information from unstructured and semi-structured human-created documents reliably, for use by other decision-support or automation tools. It also involves approaches to displaying this information in an easily digestible manner. Many military documents have some degree of structure between different instances of a document type. Extracting that information in a standard format is useful, but individual instances also contain significant variation that makes this task difficult. To extract machine-usable data requires identifying standard components present, however they are labelled and wherever they are present in the document, and to build a flexible representation that provides some structure, while nevertheless retaining the ability to represent unexpected or unknown content. In addition, the challenge is to present the information extracted in a more easily digested manner – for instance displaying key or summary information, re-arranging information (such as in a tree structure) or by providing additional meta-information to support information discovery and understanding.

8. Goal-based agent learning and reasoning in uncertain environments

In defence contexts agents must act and adapt to uncertain environments and changing circumstances. How are agents able to continually learn and adapt in the face of these challenges to improve understanding of sensing and commands? A general solution to this problem is one of the grand challenges of AI. What research, including new or innovative frameworks, exist to begin to tackle this problem? Exploration of ideas in contextual reasoning of a single agent or multiple agents communicating as a team are of interest. This could be demonstrated in a simulation environment with changing dynamics or through some new theoretical frameworks.

9. Decision making in uncertain, dynamic and complex information environments

Information-centric environments are challenging because they are fast-paced and high-dimensional; observations are not always reliable; and these environments consist of many interacting / coupled components. We seek to understand the technical issues with distributed machine-based decision making in these environments. In particular we are interested in how this decision making would differ between offline planning versus online execution, and how would one design an intelligent agent architecture to thrive in this space. A key consideration is how to realise a decision making paradigm that is scalable and amenable to contexts with urgent decision time frames, lean computational capacity and intermittent connectivity. Finally, responders are invited to demonstrate the superiority of technological solutions relative to alternatives.

10. Strategy formulation in open adversarial environments

Strategy formation in adversarial environments featuring fundamental uncertainty is challenging for humans and presently largely beyond the capabilities of machines, which rely on strong assumptions of (statistical) predictability in order to converge to fixed solutions effectively predetermined in the environment. We are interested in research that can situate existing AI or develop new AI methods capable of formulating effective strategy in open adversarial environments. These environments will be represented by adversarial games in which the governing rules can be altered by actions of players, possibly without direct knowledge of opponents. Players may also have hidden asymmetric goals and these goals may also be changeable. We would assist in selecting a game representation for exploring dynamic strategy development for machines.