



Australian Government
Department of Defence
Science and Technology

Call for Pilot Project Proposals

Artificial Intelligence for Decision Making 2022 Initiative

Challenges

The Artificial Intelligence for Decision Making 2022 Initiative is a national endeavour from the Defence Science and Technology Group (DSTG) and the Office of National Intelligence (ONI), and in collaboration with the Defence AI Centre (DAIC) to fund up to 50 pilot project proposals that will contribute to critically important Artificial Intelligence (AI) and Machine Learning (ML) technology.

The Initiative's call for proposals is "challenge based", with respondents expected to propose pilot projects that address one of 30 Challenges.

Below is a list of challenges that individuals are invited to address.

1. How can graph neural networks be applied for causal modelling?

Graph neural networks are gaining popularity for the modelling of systems of interacting elements or where one can describe a system with topological structure. It would appear that graph neural networks would be well suited to modelling causality in such systems. We seek to understand and advance the state of the art with respect to performing causal modelling with graph neural networks.

2. Musical scoring as a language of synchronised action

A musical score is a language of synchronised action. We seek to develop a formalism/language for describing synchronised and coordinated actions by teams of actors. This formalism needs to take into account the uncertainties of the environment which may be adversarial in nature. We presume that actors can communicate via limited bandwidth subject to delay and yet they are cooperating to achieve a common goal. For example, a squad of soccer players work together to kick goals and defend against their opponents from scoring goals. How can this be described using a formalism for synchronised and coordinated action?

3. What can information theory tell us about generalisation in data driven models?

Generalisation of data-driven models, such as those created with deep learning, is a serious concern. The challenge is to be able to identify the invariant nature of the underlying objects of interest in spite of the distortion imposed by either environment or imperfect sensors. Information theory

describes the capacity of channels to convey information reliably, and coding theory describes approaches for source and channel codes to approach those capacity limits. What can information theory tell us about generalisation in data driven models? How can we use information theory to infer the capacity of the channel between the object and our reasoning system? Noting that we cannot control the transmitter, how do we design our perceptual front-ends to maximise the information that we can gather from the environment about objects of interest?

4. Understanding and exploring adversarial machine learning attacks and defences in a cyber security context

Recent work has shown that the use of Machine Learning (ML) can introduce additional vulnerabilities into a system, arising from weaknesses inherent in the ML algorithms themselves. It has been demonstrated that human-imperceptible perturbations, or more generally perturbations that do not affect the semantics of the data, can lead to, for example, incorrectly classifying perturbed data samples with a high degree of confidence. The existence of such adversarial data samples results in an increase in the available attack surface of a system utilising ML for decision making, impacting the effectiveness of and confidence in deployed ML capability. Adversarial Machine Learning (AML) is concerned with understanding this threat to capability presented by adversarial data.

Adversarial attacks against ML systems in the computer vision domain have attracted the bulk of the attention to date. However, ML in the Cyber Security domain (e.g., detection of malware, phishing, intrusion etc. as well as autonomous cyber network attack/defence) presents a critical target, with extant active adversaries having clear targeted goals. Adversarial attacks in the Cyber Security remain to be fully explored, or even yet clearly defined. It is obvious that the input data in Cyber Security have fundamentally different properties compared to image samples in computer vision. Therefore, adversarial attacks in Cyber Security and computer vision may be fundamentally different. Hence the research question to address is how can the concepts, tools, and techniques of adversarial attacks on images and their mitigations be best transferred to attacks and corresponding defences against Machine Learning systems in the Cyber Security domain?

5. AI-Enabled Argument Mapping for Decision Quality

High quality decision making depends on critical thinking, situation understanding and effective communication. All three can be improved by argument mapping. Argument mapping is “visually depicting the structure of reasoning” using a “diagram, with nodes corresponding to propositions and links to inferential relationships.” [1]

The use of argument mapping in intensive courses has been shown to improve the quality of critical thinking, as measured in standard tests, by 0.75 standard deviations [2]. Mapping can uncover gaps and weaknesses in reasoning and enables more effective communication than prose or bullet point slides.

The obstacle to deploying argument mapping to improve Defence decision quality is that it consumes a great deal of effort from trained argument analysts [3]. Argument mining seeks to extract argument maps from natural language text [4], eliminating most of the cost.

Future AI technology could provide machine team members that facilitate deliberation, guiding the construction of strong arguments, and challenging proposals to uncover their weaknesses.

Possible applications of argument technology in ADF decision making include the following:

- force structure planning
- operational planning and assessment
- design of operation rooms and combat systems
- staff officer training
- intelligence analysis
- international relations analysis

DSTG seeks proposals to demonstrate state-of-the art machine assisted argument mapping in settings comparable to some of these defence decision contexts. We also want to identify gaps in the current technology, and emerging work that could close them.

References:

[1] Tim Van Gelder, Using Argument Mapping to Improve Critical Thinking Skills, in The Palgrave Handbook of Critical Thinking in Higher Education, 2015

[2] Eva van der Brugge, The use of argument mapping in improving critical thinking, PhD thesis Uni of Melbourne, 2018 <https://minerva-access.unimelb.edu.au/handle/11343/214519>

[3] Floris Bex, John Lawrence, Mark Snaith, Chris Reed, Implementing the Argument Web, Communications of the ACM, 2013

[4] John Lawrence, Chris Reed, Argument Mining: A Survey, Computational Linguistics, 2019

6. Data usage and sharing in contested or denied environments

The Navy is required to operate in contested or denied environments. These environments are characterised by the inability to use standard communications mechanisms due to the unavailability or compromise of the spectrum. At the same time there has been a proliferation of data intensive systems in Navy units. To enhance mission effectiveness this data should (ideally) be transmissible or available as and when required between different Navy assets, i.e., ship to ship, ship to shore, and autonomous system to ship. The Navy would benefit from the use of artificial intelligence to substantially reduce, obfuscate, speed up or otherwise streamline the sharing of data in the characterised environments.

7. Classification of code modifications to differentiate security fixes and learn vulnerable code patterns in compiled software

A common process for tracking and understanding software vulnerabilities is to reverse-engineer patches made available for applications. These patches inherently identify the location of the vulnerable code and the form of the underlying flaw. It therefore may be possible to use the details of these changes and the associated patch notes to train AI models of both the code to which the patch is applied and the patch itself. In the former case, this could be used to detect constructs in code that are similar to commonly fixed flaws, which may need more in-depth assessment. In the latter case, this could allow for the discernment of which of the applied changes create the fix. This

latter case is particularly relevant as not all application updates make individual changes to the application's function and may not be distributed with detail on if and what flaws are fixed.

Techniques for this research may operate by either comparing two releases of an application or the changes created by applying patches. This work explicitly targets compiled applications, as this is the predominant form in which applications and their updates are distributed.

8. A Wireless communication detection problem for Graphical Neural Networks, Graph learning and Generative Graph methods in Spatial Networks

Graph Neural Networks (GNNs) are a relatively new deep learning methods and have gained immense popularity in the research space in the last few years. GNNs can be applied to graphs to conduct node-level, graph-level, and edge-level predictions. One type of GNN - Graph Learning Networks (GLN) are a simple yet effective process to learn node embedding and structure prediction functions.

Low probability detection (LPD) communication has recently emerged as a new transmission technology to address privacy and security in wireless networks. Hiding wireless transmissions may also be explicitly desired by government and military bodies. The LPD problem is also analogous to many other practical and natural instances where reducing the detectability of communication networks is advantageous. An abstract model of the LPD problem can be embedded in a Euclidean space (both 2D and 3D versions).

In these models, each node must tune its transmit power in order to trade off radio network connectivity (and communications bandwidth) versus detectability. Varying transmit power levels will induce different graph structures that represent network connectivity and regions of detectability (footprint).

Extensions to the abstract models of the LPD problem involve introducing dynamic properties whereby network nodes and adversary locations may change, be added, or subtracted as well as introducing various footprint topologies for specific connections. These extended models may also include distributed decision-making concepts and imperfect knowledge scenarios.

The initial aim is to leverage GNNs and potentially GLNs to find efficient models that can aid in finding robust solutions to transmit power level selection that affects both connection and detection aspects of the LPD problem.

A further aim is that these solutions then can be used to inform extended abstract models for the LPD problem in a timely and computationally effective manner. This enables higher level abstraction aspects of the LPD problem, such as tactical and game theoretic approaches to be formulated and used in simulations. The GNNs used may also be repurposed to support such approaches.

The undertaking would involve integrating aspects or developing software code for a simple abstracted model and possibly the use of Python and/or Deep learning frameworks (e.g., Keras, TensorFlow) for the machine learning components. Such a challenge may generate heuristics and techniques in the use of LPD technologies.

9. Applying machine learning techniques to games on graphs for the detection and concealment of spatially defined communication networks

Providing strategies for defending communication networks such as distributed sensor systems, covert radio networks or even human voice often involves reducing detection of presence. This is a problem similar to those addressed by the low probability detection (LPD) technologies that aim to prevent an adversary being alerted to the presence of such networks. Such an adversary may either be deliberately or inadvertently seeking to find such a covert network.

Maker breaker games are a subclass of positional games in combinatorial game theory and can also be played on graphs. The detection and concealment of a connected radio network can be modelled as such a game on a graph, given both maker and breaker (adversary) have knowledge of the graph. The actions of each player are determined by geometric properties of the network in a Euclidean space (plane). Each game can be formulated by randomly generated spatial arrangements for a communication network with varying numbers of nodes. The breaker is trying to locate itself to disable as many links within the communication network as possible and cause nodes to be isolated.

Bayesian games can be used to alter and extend the maker breaker game such that aspects of imperfect games can be incorporated into the setting. This for example allows for an adversary to only have partial observability of the network.

The aim of this challenge is to involve varying types of machine learning techniques to elicit robust tactics in both the perfect and imperfect versions of these games.

The undertaking would involve generating a simple simulation of an abstracted game for various network settings and the use of python and/or deep learning frameworks (e.g., Keras, TensorFlow) code to run and analyse games. Such a challenge may provide insights into the dynamics and execution of LPD technologies.

10. Agent-based trust management

The concept of trust is essential to the secure operation of autonomous, distributed software systems. The multi-agent systems community in particular has developed a rich theory of trust, encompassing dozens of trust models. Trust has also been used in the context of authorization control ever since the pioneering work of Matt Blaze in the 1990's. Is it possible to combine some subset of these agent-based and cyber-security models to support trusted interactions between heterogeneous, adaptive agents? Which aspects of these models would be invariant to system evolution, and how could variability be captured?

11. Resilient interaction protocols

Resilience is an important property of autonomous, distributed software systems operating in highly contested environments. This property can be realised across different aspects and levels of abstraction of such systems, including agent interaction and organisation mechanisms, which in turn rely on various interaction protocols. Is it possible to develop a general model of resilience for multi-agent interaction protocols (e.g., negotiation and others)? How could this model be instantiated or concretely realised when developing new protocols?

12. Neurosymbolic BDI agents

BDI agent architectures have been thoroughly studied and widely adopted in agent-based systems. Nevertheless, the use of these architectures in practice often entails at least two deficiencies: resorting to fixed/pre-computed plans for goal achievement, and limited learning capabilities. We are investigating neurosymbolic cognition in the context of software autonomy and are interested in extending the BDI architecture to incorporate learning as well as dynamic decision-making. What options are available for such extensions (e.g., can features from cognitive architectures be leveraged?), and how would they affect an agent's performance? Which distributed decision-making approaches would these architectures be compatible with?

13. Hybrid and combined reasoning and decision-making

Recent trends in AI have shown the benefits of combining different decision-making approaches. One example is the combination of reinforcement learning with classical planning, where planning can guide the learning process as a kind of heuristic to achieve faster policy convergence. Another example is causal reinforcement learning, which also has a distributed variant: causal multi-agent RL (CMARL). We are interested in using such combined and hybrid decision-making approaches – especially combinations of multi-agent RL (MARL) and multi-agent planning (MAP), as well as causal RL and CMARL – and are seeking to understand their utility and trade-offs in the context of distributed autonomous software systems tasked with resource management as well as self-management.

14. Abstract cyber control

At present, there are several openly available environments or “labs” which allow testing and evaluation of different cyber security strategies. Is it possible to abstract away some or all characteristics of the cyber-attack/defend processes within these environments into a single, abstract game environment that still maintains relevance to real-world scenarios?

15. Modelling objectives for multi-agent systems

In the goal-driven autonomy paradigm, an autonomous system is associated with a set of objectives that can be reasoned over and modified during operation. How can these objectives be modelled? How can system-wide objectives be automatically decomposed and mapped across different abstraction levels and agent organisation structures, such that some subsets of objectives are directly usable in the reasoning and decision-making processes of individual agents?

16. Optimising human-algorithm teaming for facial recognition

Human-algorithm teaming is increasingly important in modern activities. One area that has a history of humans and algorithms working together is facial recognition. Facial recognition often involves a human submitting an image or images of a person-of-interest to a facial recognition system to determine identity. These images are compared to those stored in a database or watchlist, and the facial recognition algorithm returns the images it considers are most similar to the person-of-interest

(the candidate list). It is then up to the human decision maker to look through the candidate list and determine if any images match the person-of-interest. As the technology advances at a rapid pace, updated techniques, and methods to optimise the human-algorithm team need to be explored. This research challenge is for you to propose and test a method that may contribute to the optimisation of human-algorithm teaming for facial recognition.

Examples include but are not limited to:

- fusion methods of the human and algorithm
- the type of information presented to the human by the algorithm
- changes to how the human and algorithm work together
- screen display designs
- building trust in the technology

17. Can Electric Network Frequency be used for video geo-localisation and timestamping?

In recent years, there has been a proliferation of multimedia content connected to criminal activities, such as terrorism propaganda, cyber extortion, and even child sexual exploitation online. Metadata linked to these videos may not be available or trusted. However, geo-location and timestamping information can help to confirm where and when the video was captured, providing new investigative leads. One potential approach is to use Electric Network Frequency (ENF). ENF is the frequency at which the alternate current from the power grid fluctuates. Under certain conditions, these fluctuations can leak within audio recordings. ENF signal also has a subtle influence on the artificial light emitted by sources connected to the electric grid. This information can be captured by camera sensors producing a video, and potentially be used in digital forensic investigations to geo-locate and timestamp the video recording. Recently, the use of video to extract the ENF has been proposed to complement the audio-based approach.

Your task will be to use AI to extract ENF from video to demonstrate the strengths and limitations of ENF for video geo-localisation and timestamping.

18. Novel AI methods to support post-mortem identification

DNA, dental, and fingerprint are common methods used for formal post-mortem identification. But other methods may be beneficial to narrow down these sometimes costly and time-consuming processes. The aim of this project is to explore the feasibility of a novel AI method(s) that could aid in post-mortem identification when these traditional methods are not possible. Alternatively, these novel AI methods could be used to help narrow down potential leads before a formal identification procedure.

19. Extracting multiple sources of information from videos for identification

Identifying people in videos depicting terrorist acts, spreading misinformation, or conducting criminal activities can be difficult, particularly when the perpetrator's face is not visible. Often the same person may be depicted across multiple videos and finding information within these videos that links them together can be a considerably time-consuming manual process for investigators.

The aim of this work is to demonstrate the use of AI approaches to extract and fuse multiple sources of information from videos that can help to find the same person(s) across videos, to ultimately aid in their identification.

Examples of potential sources of information to extract and fuse include (but are not limited to):

- logos (from clothing, food items, flags etc)
- voice recognition
- tattoo recognition
- age estimation

20. Using AI with unmanned aerial systems to locate victims of a mass disaster event

Locating victims of a mass disaster can be time critical. Various sensors and AI approaches may help when attached to unmanned aerial systems (UAS). UAS are inexpensive and can be rapidly deployed to reduce the response time and risks to the search and rescue teams. The aim of this work is to advance the timely detection and location of victims of a mass disaster event. Novel algorithms to allow for automatic image analysis and faster detection capabilities should be considered. Methods should be able to be deployed on a UAS and the work should explore methods for detecting and distinguishing between living and deceased victims so that appropriate triage can occur.

21. The TTCP CAGE Autonomous Cyber Defence Challenge

The Autonomous Cyber Operations Discipline has developed CybORG (the Cyber Operations Research Gym), a framework for testing and training Autonomous Cyber Defence agents in simulated scenarios. The TTCP CAGE Challenge is an open challenge for researchers to develop defensive agents for CybORG that work effectively in a defined cyber scenario.

For this research project, the selected team will be competing within the next version of the challenge. This will require them to develop an intelligent blue (i.e., defensive) agent that can select actions to defend against a red agent moving laterally through their network. The blue agent receives a regular monitoring feed from the network, and can then choose to analyse systems in detail, remove malware if detected, restore systems from backup, and start decoy services to trick the red agent. The red agent can scan the network, scan systems for vulnerabilities, exploit these vulnerabilities to gain access to systems and then escalate their privilege. The blue agent is scored both on the degree to which it prevents the red agent establishing control of the network, and on the effect blue's actions have on the availability of the system to normal users.

Our expectation from the selected team is that they investigate and implement an innovative learning approach to the challenge, implement it, and then refine it through repeated testing. The deliverables are an agent that is effective at the challenge, together with a report on the learning methods used by this agent to choose blue actions.

22. DeepfakeOps: towards continuous integration of Deepfake technology into an evolving detection benchmark dataset

Deepfakes are AI manipulations of vision, sound and language that can be used to fool humans, with potential implications for Defence analysts and decision makers. In an Information Warfare context, Deepfakes can be weaponized to degrade situational awareness, enhance cyberattacks that gather intelligence, and spread propaganda or disinformation. While a wide range of automated Deepfake detection methods are actively being developed, the field is arguably still maturing, with some detectors that achieve state of the art results on existing large-scale benchmarks unable to identify new types of manipulation or perform as reliably on new datasets. Moreover, there is a sense of an emerging arms race between Deepfake generation and detection techniques, where the former are expected to remain one step ahead. At the same time, the AI vision, sound, and language research communities remain separate, despite the clear need for and potential benefits of multi-modal approaches to discovering misinformation. In this context, a key question for Defence is how best to track the rapidly moving target that is Deepfake technology in order to counter potential new threats quickly and at scale?

One possible solution that could be explored is to design and prototype a framework for continuous integration of new open source and/or commercial Deepfake tools that will allow Defence to test and evaluate countermeasures in a standardized manner. This could involve, but is not limited to:

Developing an automated integration pipeline allowing new Deepfake methods to be installed and run across the same target multimedia as earlier techniques used in the benchmark

Design of protocols for incorporating newly generated content into the existing dataset

Support flexible extensions beyond video, to still imagery, speech and text for future test and evaluation of multi-modal Deepfake detectors.

23. Decision Making for Network Security

AI decision making utilises data trends to develop accurate predictions and decisions for a given problem. A subset of applications includes using AI decision making for network security analytics, monitoring, and identifying responses. The use of AI to perform these functions allows improved security posture and effective responses. Areas of relevance include; the capacity to differentiate between legitimate and illegitimate traffic, the capability to monitor complex sensor networks and identify points of concern, the detection of actions that may indicate malicious activity, how adaptable models can be to a widely variable environment, and the ability to mitigate impacts from malicious actions.

What capability exists in these, and related, areas? What technology enables this functionality or would be required in future? What opportunities does this technology introduce? What risks may be presented? What predictions can be made for future applications?

24. Domain adaptation for neural speech translation

Speech translation is the process of translating from speech in a source language into text in a target language. Current approaches may model the speech translation task as either a cascade system of automated speech recognition (ASR) for transcription to text followed by machine translation (MT), or a direct speech translation (ST) model for source speech to target text without relying on

intermediate text generation. Recent studies comparing cascade to direct speech translation attest that the gap between the direct and cascade paradigms is now closed under data comparable conditions, and that subtle differences observed in their behaviour are not sufficient for humans neither to distinguish nor prefer one over the other.

However, the effectiveness of approaches in the speech translation task may be severely constrained by the availability of training corpus or corpora, especially for settings where in-domain training data is scarce or non-existent. If data is available in more readily available in different but related domains, a variety of techniques and procedures have been explored in domain adaptation and relating to single or multi-domain differences in speech translation areas such as provenance, topic, genre, dialect, or accent.

This challenge will seek to contribute new developments or understanding in speech translation for a low-resource domain setting. Some of the avenues for contributions in exploring this research area might include:

- How can an optimal approach be identified for modelling a low-resource domain speech translation task, to best leverage related resources for multi-domain adaptation?
- How to estimate or target efforts in developing in-domain corpora for greatest improvement on speech translation performance, where development capacity is tightly constrained?
- What new developments can be made in domain adaptation and/or speech translation approaches?
- What approaches in domain adaptation or speech translation can contribute to robustness and performance on the target in-domain speech translation task?

25. Translating written English to and from propositional logic

The Sensemaking team generates resilient situational insights by translating large volumes of uncertain data from disparate sources into concise and compelling findings enabling decision superiority.

Where a user's sense making needs are well defined and can be anticipated, it is possible for a programmer to implement the necessary functionality as part of a data processing capability. However, it is not appropriate to assume that this is practical given the diverse nature of those needs, the multifaceted nature of the data with which the sense maker is working, and the fact that knowledge discovery shapes the user's behaviour. Consequently, it is appropriate that the machine permits the user to pose unanticipated queries and respond in an appropriate manner.

The objective of this project is to survey the state-of-the-art in translation between human written English and propositional logic to support flexible and complex human / machine interactions.

The machine-understandable query language needs to capture uncertainty, and the translation of the machine's output back to English should accommodate both uncertainty and explanations regarding how the solution was arrived at.

There are many aspects in the natural language to machine language correspondence that need careful consideration. For instance, how does the tense of a natural language query translate appropriately to the machine language query.

Example questions to help assess such translation / sensemaking systems might be:

- What happens when a user has problems making their question understood by the machine? How robust is it to user inexperience or missing data?
- Are there tweaks that could be applied to improve the translator's performance, either by a programmer or through automated learning from examples?
- Is the software well supported, either commercially or by the open-source community?

26. Identification of significant field types and boundaries in unfamiliar network protocols

Information traversing cyber networks are encoded using network protocols that define a common language through which networked components of a system communicate. This encoded data is a valuable source of information for many defensive cyber applications, particularly for system monitoring. However, in many cases, documentation for these protocols may be unavailable or incomplete, which facilitates the need for reverse-engineering based only on examples from network packet captures. This is a time-consuming endeavour that requires a great deal of skill and experience.

This research challenge focusses on a small part of an approach to automated protocol reverse-engineering. We pose the problem of identifying the data type (e.g., integer, floating point, string) of particular fields in a network protocol based on the byte values observed in corresponding samples. As a secondary challenge, we pose the problem of identifying the location of discrete fields of a network protocol given prior examples.

A successful answer to this challenge will demonstrate an approach to automatically classify the data type of a field, and/or automatically identify the edges of fields in a network protocol. Ideally, these techniques will be evaluated across a variety of network protocols, to demonstrate a level of robustness and protocol-agnosticism.

27. Application of AI to the liveness challenge in identity verification

With the increasing use of online and remote identity verification applications for a variety of purposes (including banking, government services, and remote access control etc.), it is critical that these applications can reliably identify presenting persons. Clearly, however, verifying an identity and providing access to services, facilities or finances is highly risky without a means to ensure that the person presenting is indeed the live bona fide individual rather than a fraudster attempting to spoof the system.

Leading facial recognition algorithms are becoming very reliable in recognising individuals and matching faces. But while many also include liveness detection components designed to combat spoofing, the reliability of performance of these aspects can be much lower by comparison, which becomes a system vulnerability.

Spoofing versus spoof detection has become a rapidly moving evolutionary arms race, and while many applications claim they can detect liveness with close to 100% accuracy, the speed of evolution of spoofing attacks and materials means that these performance guarantees may not last more than a few months at best.

One of the avenues that has yet to be explored in this secure identity arena is the use of AI. Some of the issues relating to this area include (but are not limited to):

- Can AI be applied to the reliable verification of liveness for identity verification/authentication applications?
- Can AI perform better than the current algorithms on the market?
- Is AI efficient and cost effective?
- Is AI vulnerable to spoofing? And if so, what can be done to make AI more robust?

28. Cross-lingual narrative summarisation to enhance situational awareness and decision making

Open-source data has long been recognised as an important source of information which complements classified intelligence to enhance Defence strategic/ operational/tactical awareness and operations.

However due to its large volume, multi-lingual, and heterogeneous nature (e.g., news articles, opinion pieces, policy documents, and technical papers, among others), effective exploitation of this type of data is a significantly challenging task. As a result, automatic text summarisation, the process by which a document's salient content is automatically distilled into a succinct and concise form, has become a necessary tool to reduce cognitive overload.

This project seeks proposals that investigate and enhance state-of-the-art techniques such as mT5 [1], XSum [2] and XL-Sum [3] toward the automated generation of narrative text summarisation in cross-lingual setting: the ability to quickly process and understand a large collection of textual documents and summarise the most salient topics and events in the form of narratives or storylines [4]. This could assist human analysts with obtaining a coherent picture of (potentially evolving) situations of interest, necessary for conducting further analysis and responses.

References:

- [1] Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., Barua, A. and Raffel, C., 2020. mT5: A massively multilingual pre-trained text-to-text transformer. NAACL 2021.
- [2] Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. EMNLP 2018.
- [3] Tahmid Hasan, Abhik Bhattacharjee, Md Saiful Islam, Kazi Samin, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, Rifat Shahriyar, 2021. XL-Sum: Large-Scale Multilingual Abstractive Summarization for 44 Languages. ACL-Findings 2021.
- [4] Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-And-Write: Towards Better Automatic Storytelling. AAAI 2019.

29. Online and Adaptive Learning with Irregularly Sampled Data in Uncertain and Non-stationary Environments

Military operations may need to resolve rapidly evolving situations where adversaries are adapting their tactics, techniques and procedures and the behaviour of populations is changing.

AI, ML and deep learning offer opportunities to better exploit information to improve understanding, decision-making and tempo toward more agile decision making.

However, deep learning methods often require massive amounts of reliable, trusted data collected in stationary environments - inadequate to deal with adversarial and non-stationary data and environments, particularly at the edge where compute, communications and labelled data is limited and arrives incrementally in real-time.

We seek proposals that investigate and enhance state-of-the-art techniques toward endowing a deep learning model an ability to evolve in uncertain, rapidly changing contexts. Specifically, we are interested in research that can address the time-series prediction task in an online and adaptive learning manner [1] [2]. In contrast to typical offline learning (where all training data must be available at the time of model training), the online learning paradigm continuously updates a model during operation whenever circumstances change and/or more training data arrives. In addition, the framework should also be able to handle irregularly sampled time-series data that is more likely to be obtained in uncertain, changing and non-stationary environments.

References:

[1] Abushaqra, F.M., Xue, H., Ren, Y. and Salim, F.D., 2021, December. PIETS: Parallelised Irregularity Encoders for Forecasting with Heterogeneous Time-Series. In 2021 IEEE International Conference on Data Mining (ICDM) (pp. 976-981). IEEE.

[2] Wang, Z., Jiang, R., Xue, H., Salim, F.D., Song, X. and Shibasaki, R., 2021. Event-Aware Multimodal Mobility Nowcasting. arXiv preprint arXiv:2112.08443.

30. Adaptive Modelling Across Domains and Tasks without Labelled Data

Modern warfare is no longer constrained to military operations in the physical domains (land, air, sea and space), but has expanded to include information and influence operations in non-physical domains such as information/cyber and cognition.

Military/information/influence operations enhanced with AI hold promise to improve situational awareness, increase the speed and scale of decision making and better conduct distributed and complex operations. However, it is not scalable nor effective to curate a great amount of labelled data and develop a bespoke AI solution for each single task and mission we plan and conduct within each domain and distinct environment.

Self-supervised learning [1] has been identified by DARPA as an important technology in the third wave of AI [5] toward extracting 'the dark matter of intelligence' [6] without relying on a large amount of labelled data as in supervised learning.

We seek proposals that investigate and enhance state-of-the-art techniques toward the development of deep learning models that can be quickly adapted and applied to new domains, environments and types of data which is often unlabelled 'in the wild'.

In particular, this challenge seeks to explore an adaptive modelling approach that can be used to learn representations of various types of time-series data (e.g., audio and sensor data) [2] [4] in a self-supervised learning fashion without requiring well-annotated data [3].

A key metric for success is that the framework can adapt to different environments, across tasks, across domains.

References:

- [1] Chen, T., Kornblith, S., Norouzi, M. and Hinton, G., 2020, November. A simple framework for contrastive learning of visual representations. In International conference on machine learning (pp. 1597-1607). PMLR.
- [2] Xue, H. and Salim, F.D., 2021, August. Exploring self-supervised representation ensembles for covid-19 cough classification. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (pp. 1944-1952).
- [3] Deldari, S., Smith, D.V., Xue, H. and Salim, F.D., 2021, April. Time series change point detection with self-supervised contrastive predictive coding. In Proceedings of the Web Conference 2021 (pp. 3124-3135).
- [4] Saeed, A., Salim, F.D., Ozcelebi, T. and Lukkien, J., 2020. Federated self-supervised learning of multisensor representations for embedded intelligence. IEEE Internet of Things Journal, 8(2), pp.1030-1040.
- [5] Daws, R. 2018. DARPA introduces "third wave of Artificial Intelligence". AI News. <https://artificialintelligence-news.com/2018/09/28/darpa-third-wave-artificial-intelligence/>
- [6] MetaAI, 2021. Self-supervised learning: The dark matter of Intelligence. AI Facebook. <https://ai.facebook.com/blog/self-supervised-learning-the-dark-matter-of-intelligence/>