

AI for Decision Making challenges

1. Facial Recognition Bias Profiling

The goal of this task is to identify when potential facial recognition models are unfairly biased. We are primarily concerned about bias regarding ethnicity, gender and age, although we welcome the inclusion of other biases. We would like a method to evaluate models and get indications as to the level of bias present. One potential approach is to gather a balanced, labelled set of data to run a model across. The data set wouldn't have to be big enough to re-train the model to account for biases.

2. Facial Verification Explainability Measures

For the facial verification task – are two given faces the same person? – build an interpretable explanation mechanism as to what facial features or details either supported or refuted the outcome. The ideal proposal will include comparisons across the two images, rather than saliency maps on each individual image. As a starting point, the concept demonstrated in arXiv:2003.05383 is promising, but lacks a useable public implementation. We would also be interested in other ways of visualising these explanations, with the end goal of making the explanations as simple to understand as possible.

3. Deceitful/Persuasive Writing Detection

We are looking for a method to detect when an author is intentionally exaggerating or being deceitful amongst a corpus of documents they have written. That is, while the authors are technically the same, can we detect where an author's writing style may have changed (using lexical changes associated with deceit, exaggeration or similar) for documents making grand claims? As a starting point, a modified version of the 'General Imposters' technique could be an interesting avenue to explore.

4. Distribution-Aware Generalisation for Imbalanced Learning

When handling imbalanced target variables, eg. there are far fewer positives than negatives, it is often the case that the over- or under-sampling will modify the distribution(s) of the input variables. We are looking for a method to regularise a model and minimise the disruption to the input distributions while better handling the imbalanced target variables. An extension to this would be applications for the online learning domain, where we have an incoming stream of imbalanced data and frequent (or continuous) retraining of a model.

5. Human Factors in Explainable AI

We are after a set of user experience designs and principles that can adequately communicate uncertainty. Machine learning models inherently come with some uncertainty, and deep learning models in particular can be difficult to interpret. Analytic systems that don't expose their inherent uncertainty can result in a loss of user confidence in the system. We are therefore looking for UX techniques that can mitigate this loss of trust while offering useful ways to manage uncertainty in the analytic products, or the data itself.

6. Input Feature Obfuscation

We are after a method of obfuscating the relative importance of input variables to a ML model. The goal here is to mask the contributions of input features to an adversary who can supply inputs and observe outputs from the model. This could be done either by playing with the model weights at inference time, adjusting the prediction value within reasonable bounds, or by obfuscating the variables before they are fed to the model. Other approaches are, of course, welcome.

7. Low-bias Age Prediction

We are after a pre-trained model to predict someone's age from a photo of their face, while minimising ethnic, age and gender biases. We are comfortable predicting either a specific age (eg. 37 years old) or predicting age categories (eg. 35-40 years old). The expected output of this problem would be a working age prediction model, together with evidence that these biases have been addressed in some way.

8. Ensembling Single-Class Classifiers

We are after techniques for ensembling single-class classifiers to improve outlier detection. Ensembling can be done by either creating a suite of models with different hyperparameters or by partitioning the input feature space and training models over each partition. In either case we need to aggregate the predictions of many models, and in the partitioning case we need to define useful partitions over the input feature space. We are looking for a literature review followed by an open-source implementation of any state-of-the-art or innovative techniques, as well as the results of experiments on publicly available datasets.

9. Scoping the shopping Race

You plan to drive from your home to a supermarket in your area in order to purchase a much in demand vital item, and you have several choices to make:

- The time of your effort to make a purchase, which can be either morning or afternoon.
- The particular supermarket you will visit, knowing that you may only pick one.
- Which of the three possible routes you might take to get to your chosen supermarket, assuming that you must pick your route in advance. All routes have become congested recently, and becoming jammed in traffic is a failure to obtain the vital item.

You know that two of the supermarkets will be restocked with the vital item in either the morning or afternoon, but will sell out rapidly since supply on this item is only 15% of demand (as measured by the supermarket managers). You have enough for only two days so your need to succeed is high. The goal is therefore to compete successfully with other shoppers seeking the item by maximising your likelihood of obtaining it.

You also know that most people competing for the item will be using the new Google/Waze service, which utilises state-of-the-art game theory and machine learning. Yet you are aware that this service is poor at handling the unexpected situation in your town during this unprecedented time: route planning involves unknowns and surprise. While traffic patterns and supermarket shelf restocking are able to be statistically learned by the Waze, availability

of the vital item is apparently random, other than what the supermarkets manage to obtain appears directly on their shelves when they open for trade.

You have access to all the past Waze data: route successes and failures, plus all the records of where and how much of the item appeared on shelves, and how rapidly it sold.

You have six friends who all need this vital item. What scheme might your friends and you employ to use human learning to outsmart Google and obtain the item against your competitors? Should you act alone, or is it worth organising a large percentage of the community to follow a common plan?

The approach you take has wide potential applicability, because the problem mixes some elements that can statistically predicted with elements of soft understanding and unpredictability. Despite the uncertainty, you still have to make decisions.

10. Understanding the Limits and Biases of AI using Connect-4

The Bigger Picture: AI is good in playing games it knows. But what if something about the game changes that, while the rules are the same, AI finds itself suddenly in need to develop new strategies to play the game? What if all the strategies that AI has used before do not apply to this new situation? How does the bias in what the AI knows affect its performance in novel situations? We would like to use Connect 4 to understand (1) how would a human or an AI manage this situation? (2) What makes the new situation complex? (3) What would be an appropriate set of steps to follow to adapt to the new situation and the bias associated with that?

The Problem: Connect 4 is a simple two-player game that we almost understand it completely. Those who understand the game know how playing first gives them an advantage. There is no luck or chance (such as rolling a dice) in Connect 4. It is all about strategy except that the strategy space is limited.

Instead of having two players, let us imagine we make three players play Connect 4. The game is now very different. Let us call these players: Mady, Marley and Maryanne. If Mady starts playing yellow, Marley plays red, Maryanne plays yellow, then now Mady finds himself/herself playing red. If they continue to play this way, the strategy space for Connect 4 is now very different. Each player needs to think about winning in a very different way. They may need to set themselves to fail in one step with a particular colour to win when their next move is due in the other colour. They have to think ahead. The basic rules of the game have not changed. What changed is the number of players which impacted a single implicit rule in classic Connect 4: that a player plays with the same colour.

What is the impact of having 3 players playing connect 4? Is there a clear winning strategy? Is it still the case that the first player would have an advantage? What if two players align in a coalition that makes the game turn into a trivial game. What if we have 5 players and extend Connect 4 to Connect 8? How complex will these setups get?

Skills to solve the challenge: this challenge is designed to be tackled by anyone who sees it interesting. Based on your skill set, you could do a lot still. For example:

- **A person with interest in Games but No Skills in Computer Science or Mathematics:** we want you to document how you are thinking from the second you start thinking about this problem. Play with your friends and kids, and record every game and every move each player played. Record also the time of the move. Document how you are thinking during the game, why you made a particular move, and post the game, reflect and see what strategy you followed and what strategies the other two players followed. Do the recording in an excel sheet, timestamp it, and write your reflection in a word file. Design strategies and rules from your experience on how to win. Compare these strategies to the classic Connect 4 and identify what needed to change and the bias in using the same strategy of classic connect 4 in this novel multi-player connect 4. Document all your work in a report and name the excel sheets properly using game setup, date, time, and any other information useful for someone else to analyse the data.
- **A person with Mathematical Skills:** Contrast the state and search space of the new 3-player game with the old 2-player game. Analyse the complexity of the game. Prove, or at least analyse, if the person opening the game would still have an advantage or not. If you are familiar with NP-completeness, show the complexity of a connect L on an MxN board and study the impact of the number of players, P, on the complexity of the strategy space as a function of L, M, N and P. Which game design (parameterisation of L, M and N) would a connect L on an MxN board approaches a 3-SAT complexity? Is there a phase transition or a critical point for this complexity to emerge? Document all your analysis in a report.
- **A person with Computer Skills:** use an open source for connect 4 (see for example <https://github.com/evison/Connect4>) and write code with different strategies to play the 2-player game. Use the same code but with 3-players and conduct an experimental study to compare the scores of different strategies and their performance in the 2 and 3 player game. Increase the size of the board and change the game to connect 5, 6, and so on. Repeat the experiments for these new setups. Collect the data and curve fit. Study the bias. Study the coefficients of the functions you fit and draw conclusions on the behaviour of these coefficients against the setups you had to answer the questions raised in the problem definition. Document all your code, data and analysis in a report.

11. Regret based strategies in a search market

In a search market, a number of consumers seek to transact a single good or service from one of a number of sellers. There are two kinds of consumers. One group, the informed consumers, "shop around" - they search some or all of the sellers and select the one with the smallest price. These are the people you see looking up store catalogues, searching the internet or running around between stores. The second group of consumers are uninformed - they simply choose a seller at random and pay the asking price. These are the people who either don't care about the cost, or who are not prepared to pay (either money or effort) to shop around.

The sellers each set their prices independently. The question is how do the sellers set their prices in some optimal way?

The symmetric Nash equilibrium (one where all sellers obtain the same profit) can be computed in closed form. It has a "U-shaped" price curve. There is a concentration of posted prices at a minimum

value which can be determined, and also at the maximum willingness to pay. There are not many sellers who post prices in the middle of this range. This isn't surprising, because whilst we are used to seeing sales with 50% or more discounts, we don't see sellers offering small discounts - they just don't attract the informed consumers.

This problem is about sellers trying to learn optimal behaviour by observing the transaction process over time. The idea is that each seller reasons about the outcome of past transactions and constructs a "regret" - this is how much the seller thinks they could have improved their profit if they had posted different prices in the past. Regret can be used to update prices in such a way that the overall price setting regime converges to some optimal strategy.

This problem is significant in internet commerce which is going to be a dramatically increasing part of our economy, especially after the COVID-19 pandemic permanently changes the way people interact, including commerce. Solution methods based on this approach can be used to automate many e-commerce activities and could be incorporated into various platforms already in use, or in new e-commerce platforms which will emerge after the COVID-19 crisis has abated.

12. Newcombe's Dilemma

This thought experiment involves a paradox of prediction that forces two distinct decision-making modes, which usually operate under different circumstances, into direct conflict within a single situation. The puzzle highlights the inapplicability of simple conventions of rationality as reward maximisation in application settings involving fundamental uncertainty. The first decision-making mode is the impetus to act in order to produce a desired outcome, and the second is the impetus to act only when the action can alter outcomes.

The dilemma is that you are presented with two boxes: one open and one closed. The open box contains a thousand dollars, while the closed box contains either a million dollars or nothing. You can choose either the closed box or you can choose both boxes. The contents of the closed box has been prepared in advance by an oracle machine that almost perfectly predicts in advance what you will choose: if you choose both boxes, then it has put nothing in the closed box, while if you choose just the closed box then it has placed the million dollars in it.

The solution to the dilemma seems obvious to almost everyone; the catch is that people divide about evenly on which of the two solutions is the 'obviously' correct one. In other words, the dilemma poses two incommensurate choices that are both justifiable. The dilemma was only fairly recently resolved using game theory to show that the course of action in the game depends on the observer's beliefs – which are necessary to making a choice but not rationally justifiable within the situation – about the ability of the oracle machine to predict their actions, in much the same manner as the observation of quantum states depends on the observer.

1. The base problem is to look at an iterated version of the dilemma for AI players, and AI oracles, playing the game repeatedly. Both will have access to all past data about the outcomes of previous rounds. Under what conditions can the player form rational beliefs about the ability of the oracle to predict their actions? What power can an oracle have to accurately predict the actions of the player?

- more mathematically inclined applicants may formally analyse the iterated game;
- computer science and AI focussed applicants may systematically experiment with learning strategies in the iterated game.

Now to a more advanced extension of the puzzle: meta-Newcombe's problem, which has not been formally resolved. The predicting oracle machine may here elect to decide whether to fill the closed box after the player has made their selection, but the player is not allowed to know whether or not the oracle has chosen to fill the closed box.

There is also another oracle in the form of a meta-oracle predicting player's choices that has reliably predicted both the player's and the oracle's choices in the past. The meta-oracle now predicts that either the player will choose both boxes, and the oracle will make its decision after them, or the player will choose only box the closed box, in which case the oracle will already have made its choice.

Choosing both boxes now creates a new dilemma. If the player chooses both boxes, the oracle will not yet have chosen what it will do, and therefore a better choice would be for the player to choose only the closed box. Yet if the player chooses only the closed box, the oracle will have already made its choice, making it impossible for the any decision by the player to have any effect on the oracle's decision.

2. The second set of problems concerns this extended game with a player, oracle and meta-oracle.
 - a. mathematically oriented applicants may formally analyse the extended puzzle to explore the nature of beliefs as external choice to the game situation towards resolving the dilemma.
 - b. computer science and AI applicants may experiment with learning strategies for players, oracles and meta-oracles in an iterated form of the extended puzzle, where the extended game is played repeatedly. Players, oracles and meta-oracles will have access to all past data about actions and outcomes of previous rounds. Under what conditions can the player form rational beliefs about the ability of the oracle to predict their actions? How accurately can an oracle predict the actions of the player? How accurately can meta-oracles predict the actions of the player and of the oracle?

13. Holiday planning by AI

You and your co-workers (working in the health care system) decide to take a much-needed extended holiday once a vaccine has been developed and the quarantine restrictions have been lifted. Yet you are a diverse group with different expectations and different levels of travel experience. You get together and rate your different potential holiday destinations, assessing how well each destination meets your different holiday objectives, and the different levels of risk associated with each of the destinations. You provide both a quantitative assessment and a more important qualitative description of why you have assessed it in that way. The holiday destination scene is now very fluid, with options and risk factors changing daily. A typical human-centred approach to this problem would use discussion by a panel of experts to argue the different considerations and reach a consensus (or impasse). The challenge is to utilise AI capabilities to assess and visualise the results of your assessments to identify any potential consensus, any particular outliers, and the safest and most risky options to your co-workers, and construct an argument given any viable set of such assessments that explains these conclusions.

An example data set can be easily generated by around a dozen friends to populate a table with around a dozen entries using headings such as:

- Holiday destination,
- Level of excitement,
- Level of risk,
- Confidence in assessment,

- Reason for assessment.

14. Document Corpus Analysis

The challenge is to build up an application that takes, as input, a set of one or more keywords from a user and searches Wikipedia (or other open-source document repositories) to locate material related to those keywords. Once a set of documents/texts related to the topic are found, the application is to provide analysis in an easily-digestible manner, focused on the sentiment, range of perspectives and context of the information. The precise way that material is grouped and presented is open to the participant - but the intent is to use automated sentiment analysis or related techniques on the material in order to display helpful information for a decision maker interested in the keywords - with the focus on highlighting the range of perspectives and sentiments expressed in the source material. The challenge is to develop some novel ways of classifying and/or visualising this kind of information, to help a decision-maker better explore the range of perspectives that might be expressed in any large corpus of documents. In addition to including techniques such as sentiment analysis to explore the range of sentiments being expressed, this challenge might include representations of other information that might be useful, such as alternative meanings or definitions of keywords, or highlighting areas of ambiguity or disagreements in various source material. This task is deliberately open-ended to encourage novel ideas in this space – and could include new ways of displaying or representing source material as well as searching through and analysing it.

The overall intent is to demonstrate one or more interesting ways of searching through any large corpus of documents and displaying a subset of material with a varied set of perspectives in order to help a decision maker get a better sense of the range of perspectives, meanings and sentiments present, without necessarily reviewing all of the material.

15. Dynamic decision-making AI for penetration testing

Penetration Testing (or “pen testing” for short) aims to improve the security of a computer network through attempts to breach that security and gain access to assets of value within the network.

If a pen tester is looking for a particular file of interest in a network, they may have to traverse a number of machines through a network before they find the file of interest. The pen tester must gain access on one machine before being able to use it to discover or access other machines. The effort required to attempt to access a machine, and the likelihood of successful access, are both highly dependent on the configuration of each machine. This problem can be expressed as exploring a graph to find a particular vertex but where: (a) traversing an edge to a new vertex has a probability of failure (in which the cost of traversing the edge is incurred but the edge is not traversed), and (b) the edges connected to a vertex, their probability of failure, and their cost are only discoverable once that vertex is reached. These probabilities can also change.

Most methods for learning decision-making assume a fixed action space. For this problem, the action space changes dynamically as the pen tester discovers more systems. A simulated version of this problem that could be used for developing learning agents can be provided. A potential extension of

the problem would be to train a defensive agent to choose actions (in particular, removing edges or vertices) in order to most efficiently hamper the pen tester.

16. Graph-based analysis of network traffic for Cyber applications

Australia's reliance on IP-based communications continues to grow unabated. The consumption of online services such as emailing, chat messaging, teleconferencing, news and social media, gaming and audio/video streaming make up a significant portion of our daily routine. This has resulted in rapid growth in the volume, velocity and variety of internet communications traffic, which limits our ability to protect network traffic data and maintain Cyber security. This challenge is to build on contemporary research to prototype affiliation graph inspired analytics to enhance identification and/or discovery of suspicious network traffic.

The project involves two complementary research tasks:

1. Research and prototyping of affiliation graph analytics incorporating machine-learning algorithms to characterise the online services and application behaviours belonging to network traffic. A selection of clustering and related supervised classification techniques shall be implemented and evaluated against existing methods.
2. Creating an informative and interactive dashboard-inspired analytical framework to present results and highlight important outcomes from the above affiliation graph mining techniques.

The outcome of the project is a proof-of-concept tool which can identify suspicious network traffic.

17. Machine learned models from simulation data

The idea for this project is to evaluate how well data generated by the simulation of a model can be analysed using machine learning techniques to create a meta-model of the simulation itself. This meta-model is especially desirable if it is more efficient than the simulation in terms of time and/or computational resources.

A model is simulated over a large number of inputs and the corresponding outputs collected. Relationships between inputs and outputs can be discovered via machine learning techniques and a "meta-model" of the simulation created. This meta-model can be used in place of the simulation should an input vector not previously simulated need to be analysed. The meta-model would interpolate/extrapolate based on the data from which it had been learned. Alternatively, an "inverse" meta-model could be learned so that inputs (or ranges of inputs) can be discovered that yield desired outputs.

Some challenges to explore:

- Can the "validity" of the output from the meta-model be quantified? Once the validity falls below some defined quality level it should trigger/flag the rerun of the simulation to gather more data and update the meta-model.
- The output of some simulations will be a time series of data for each "tick" of the simulation clock. Other simulations will have one output at the end of a set number of time steps or when pre-defined end conditions are met. What challenges are faced in applying ML techniques to these different simulation types?

An epidemiology model (there are many freely available) might provide a good test case for exploring these ideas.

18. Using Machine Learning for Game and Player Analytics

Simulations generate large datasets which are ripe for analysis. Many simulations used in the operations research domain resemble either traditional games, real world sports games or e-sports games. That is, they typically involve two teams in an adversarial manner trying to out-score each other.

Many factors determine the outcomes of a game. They include the individual skills of players on the team, the skills of the team as a whole, the tactics employed as well the skills of the opposing team and environmental effects such as weather.

In this project we are interested in exploring sports data sets both from real sports results and from e-sports. The general idea is to use both supervised and unsupervised machine learning techniques to find relationships of interest in the data.

The researcher is free to use any ML data analysis techniques to explore a number of proposed datasets and is free to propose their own questions. In addition to the application and interpretation of ML techniques the project will also involve generating appropriate visualisations and interpretations of the results.

Proposed Datasets

It is recommended that the researcher start with these datasets. Once they have been exhausted, other similar sports data sets may be explored if time permits.

- FIFA 18 Player Dataset

<https://www.kaggle.com/thec03u5/fifa-18-demo-player-dataset>

This data set contains information about all the players represented in the FIFA 2018 video game. This blog post:

<https://towardsdatascience.com/exploratory-analysis-of-fifa-18-dataset-using-r-ba09aa4a2d3c>

provides some simple examples of the analysis that can be performed on this dataset.

- FIFA World Cup Results

<https://www.kaggle.com/abecklas/fifa-world-cup>

This data set contains results from all World Cup finals matches

- Football Events Data Set

<https://www.kaggle.com/secareanualin/football-events>

This dataset contains over 900,000 events from over 9,000 European football matches.

- European Soccer Data Set

<https://www.kaggle.com/hugomathien/soccer>

A data set with information on over 25,000 European league football matches.

Machine Learning Techniques

The researcher may use any standard or advanced ML technique (or combination of techniques) to address the challenge. It is expected that the project will make use of open source tools and libraries. Some example ML techniques which should be considered include regression, clustering, principal component analysis, decision trees, random forests etc.

The Challenge

The challenge for the researcher or research team is to apply ML techniques to these datasets that will extract the most useful information and insights on player, event and match result data. It is expected the researchers will pose meaningful questions and attempt to answer them with the relevant techniques with statistical rigour. Explanations of how the techniques are used and supporting visualisations should also be provided. All code developed on this project also needs to be provided.

Some example questions that the researchers may want to consider to get started may include:

- What features contribute the most to a player's salary and/or value?
- What features determine the outcome of a match?
- How can we visualise the distribution of match outcomes?
- One way of grouping/clustering players is by associating them with the position that they play in? What are the other factors that determine how we group players together?
- Can we predict how many goals a player will shoot in a season?

Success in the project will be considered based on the ability of the researcher to show creativity and to be innovative in how the data is analysed, interpreted, presented and visualised.

19. Uncertainty and Deception in Multi-agent OpenAI Gym Environments

Real-world problems have many aspects that make them difficult to solve using state-of-the-art reinforcement learning (RL) approaches, differing from games like Go and Chess due to factors such as asymmetry, partial-observability, continuous state or action spaces, teaming, multiple objectives, or uncertain rewards.

OpenAI Gym is a research toolkit that provides a standard set of environments for developing and testing RL algorithms. The challenge for the researcher is to develop one or more multi-agent OpenAI Gym environments that capture complexities of real-world problems, and to demonstrate an RL algorithm against these environments. A list of example problems is provided below. Researchers may expand these or develop their own unique games to address the challenge.

- **Hide and Seek:** One or more seekers must locate and tag a hider before they reach home base. The hider may use camouflage and distractions to avoid detection and lay false trails, while the seekers may place traps that sound an alarm when the hider passes nearby.
- **Flashlight Tag:** Two teams of players in a dark room try to tag each other until one team is victorious. Players have a torch with a narrow beam, and must balance the competing objectives of protecting themselves while finding and tagging their opponents.
- **Warehouse Patrol:** A pair of security guards must determine the most effective and efficient way to search and clear a large warehouse with an unknown number of intruders. They have limited time and the intruders may be armed and dangerous.
- **Rescue the General:** An ambulance must reach a General who has been injured on the battlefield, while preventing enemy soldiers from determining the General's whereabouts by observing the actions of the ambulance.
- **Deception:** Modify a multi-agent OpenAI Gym environment to allow a player to manipulate the observations provided to another player, and demonstrate the difference in the behaviour of an RL algorithm between the modified and control environments. For example, one player may be provided with an action that reduces the observation range of their opponent, or presents a false location.

A multi-agent interface to OpenAI Gym will be provided, and researchers are free to modify existing environments to address the challenge. The researcher may use any standard or advanced RL techniques (or combinations of techniques) to demonstrate their environments. It is expected that the project will make use of open source tools and libraries, such as the set of baseline algorithms provided by OpenAI (github.com/openai/baselines). Questions that the researchers may choose to consider include:

- How will this environment capture a challenging aspect of real-world problems?
- Which RL approaches are most appropriate to the problem?
- What aspect of the observation space could an agent logically affect for another agent?
- How does agent behaviour change when it can modify the beliefs of its opponent?

20. Automated glossary generation for effective and efficient information extraction from text data

Knowledge contained within textual data can hold key information to enable superior decision making. However, extracting right knowledge and deriving evidences required by analysts and decision makers is often dependent on creating domain-specific glossaries that are both time consuming to develop and harder to port across multiple applications. Applying AI techniques (within Natural Language Processing and broader area) can inject efficiency to this process by automating the building process for glossary. For the purpose of this problem, a glossary is a content analysis dictionary that can be referenced to extract relevant information from a corpus. It should ideally have a structured list of terms, phrases and concepts relevant to the terms arranged in some specified category.

The challenge, thus, is to:

1. Given a reasonably sized corpus and relevant list of questions, how do we automatically generate rich glossaries (or their models) that can then be applied to extract nuanced information from the corpus across multiple domains?
2. How is the hierarchy of terms/concepts structured within such glossary models?
3. How can a user visually inspect and interact with the glossary model and fine tune it if and when necessary?

Bonus task:

Can the glossary generated predict the relative coverage of information within the corpus to confidently answer a particular question or group of questions?

Contestants are free to pick up a development language of their choice. Use of credible free open source resources is permitted. The deliverable will be the algorithm and software codes together with documentation.

Contestants are encouraged to use the dataset for “COVID-19 Open Research Dataset Challenge” (<https://www.kaggle.com/allen-institute-for-ai/CORD-19-research-challenge>) as a corpus and the key task questions (<https://www.kaggle.com/allen-institute-for-ai/CORD-19-research-challenge/tasks>) as guides to demonstrate automated glossary development process and its efficacy.

21. Classification of subtle activities

A classifier is required to address the information content on both single frame digital images and short video sequences.

- The static problem consists of identifying items being carried by personnel (for instance, shovels, rakes, brooms, or weapons).
- The dynamic problem consists of recognising a range of activities being undertaken by personnel (for instance, tying a shoe lace, burying an item, walking, or running) from short video sequences.

The performance of the candidate classifiers could be judged against accuracy of classification, against different spatial resolution of imagery (stills or sequences), and against latency of the classification and the required computational resources needed to execute it. Applicants may also consider additional performance metrics.

The work should provide a vocabulary describing the features for classification of types of objects and types of activities, respectively, in general, which would then be populated with a number of examples for testing and evaluation. Algorithm development for classifiers would then utilise this vocabulary with any particular vocabulary populated with items and activities of interest as an input.

22. RTS Game Disruptor

In commercial computer games, a large amount of effort goes into testing and tuning so that the game presents the right level of challenge to match a player's skill. In multi-player games, such as Real Time Strategy (RTS) games, game balance is an important aspect, to ensure that the resources available to any one player are not unfairly superior to those available to another. Such games typically employ a range of offensive and defensive platforms with complex intransitive relationships between them, to ensure each platform has some kind of weakness and is not 'unbeatable'. A number of automated attempts have been made to automate the balancing task, for example (Wheat et al. 2015) and (Preuss et al. 2018).

The goal of this challenge is to 'break' games by giving one side an 'unfair' advantage - thus disrupting game balance. The goal is to be achieved through the design and development of an automated approach that will quantify how variables within the game can disrupt game balance.

References:

Wheat, D., Masek, M., Lam, P., Hingston, P., (2015), Dynamic Difficulty Adjustment in 2D Platformers through Agent-Based Procedural Level Generation. Proceedings of the 2015 IEEE International Conference on Systems, Man and Cybernetics (SMC 2015), 2778-2785, Piscataway, N.J., IEEE.

Preuss, M., Pfeiffer, T., Volz, V., & Pflanzl, N. (2018). Integrated Balancing of an RTS Game: Case Study and Toolbox Refinement. In 2018 IEEE Conference on Computational Intelligence and Games (CIG) (pp. 1-8). IEEE.

23. Deep Reasoning Reinforcement Learning for Cognitive Information Warfare

This challenge aims to utilise human expertise to extend the capabilities of Reinforcement Learning (RL). RL is a powerful and flexible machine intelligence technique for learning how to control the

behaviour of an agent in a system for which the system dynamics is fixed, but either parametrically unknown in advance or that exhibits high dimensionality. In these systems, traditional control approaches can break down, while deep RL (RL based on deep learning) has been demonstrated to be well suited. However, RL can also be expensive in terms of the number of samples that the approach requires to generate behaviour that is effective. One method of mitigation of this is Propositional Logic Networks (ProLoNets), which attempt to represent human expertise as decision trees. They therefore promise to offer a means of combining human subject matter expertise with machine optimisation.

Specific activities within this challenge may include:

- Review of literature relating to warm-starting reinforcement learning using human subject matter expertise, with the goal of identifying a class of algorithms of which ProLoNets is one exemplar.
- Implementation of exemplar algorithms as agents compatible with the AI Gym environment.
- Experimentation of exemplar algorithms on benchmark AI Gym RL problems, with the goal of characterising performance. Part of this effort is the identification of suitable metrics (e.g. performance as a function of computational budget, explainability or interpretability of agent behaviour), and selection of benchmarks suited to characterising performance.
- Investigate control policy complexification as a function of training time.
- Functional, statistical and theoretical analysis of exemplar algorithms with reference to metrics in order to identify deficiencies, and recommend improvements.
- For the exemplar algorithms, investigate control policy complexification as a function of training time, such as whether the complexity is bounded by implicit design choices, or can the policy achieve arbitrary levels of complexity? Can the policy escape local minima in policy solution space?
- Investigate the ability of the exemplar algorithms to generalise from training environments to environments with different initial conditions.
- Investigate methods of integrating unstructured and unlabelled data with exemplar algorithms.

Work in this topic area should aim to yield general machine intelligence algorithms, insights and recommendations that can be applied in different environments, especially online environments where the machine will act and learn as timescales too small for humans but will be responsive to changing circumstances that do evolve in human time scales.