Title: Detecting Earthquakes using AI (Data Mining, Machine Learning, Seismology)

Supervisor: Dr Azzopardi Joel

Description: Seismology refers to the study of earthquakes, and earth movement. Earthquakes are caused by seismic waves – the release of energy from geological faults (faults in the earth surface). This releases elastic energy that cause movement (or shaking) of the earth's surface. Nonetheless, not all earth's surface movement are caused by earthquakes – human activities can cause some earth surface movement as well.

Seismology involves collection of data from a network of seismographs which measure earth's movement, and a core application involves the filtering of actual earthquakes from ambient noise.

This project involves the application of AI techniques to data collected from seismographs to identify earthquakes from noise that have occurred including identifying small local earthquakes that could not be detected by traditional methods.

This project will be supervised in collaboration with the Seismic Monitoring Research Unit within the Department of Geo-sciences in the Faculty of Science (University of Malta).

Recommended Resources:
https://www.sciencedirect.com/science/article/pii/S2666544120300010

Ethical Issues:
N/A – this dissertation is envisaged to make use only of publicly available datasets

Title: Datasets in Research (Information Retrieval, Information Extraction, Data Mining)

Supervisor: Dr Azzopardi Joel

Description: One of the main challenges involving research and data-driven services is the identification of relevant datasets that can be utilised. When faced with such an issue, one typically uses search engines for this task. However, search engines do not specifically identify datasets, and researchers/developers need to go through a number of search results to identify relevant datasets. Moreover, dataset descriptions may contain minimal textual description that would help search engines index them.

A possible way to assist in dataset retrieval is to identify datasets used in similar research – i.e. to analyse publications describing similar research, and identifying what datasets have been used. The mapping of datasets to scientific publications is also useful to measure the utility of that dataset.

Kaggle in recent years, announced a competition who objective is to solve this problem – i.e. identify the datasets mentioned within similar publications (https://www.kaggle.com/c/coleridgeinitiative-show-us-the-data/data last accessed May 2023). The scope of this project is to tackle this same problem. Furthermore, this research can be used to demonstrate how the solution to this problem can help in dataset retrieval.

Recommended Resources:
Personal Computer
https://www.kaggle.com/c/coleridgeinitiative-show-us-the-data/data

Ethical Issues:
N/A – this dissertation is envisaged to make use of publicly available datasets

Title: Environmental Impact on Agricultural Output (Data Mining, Machine Learning, Agriculture)

Supervisor: Dr Azzopardi Joel

Description: One of the main challenges faced by modern society is sustainable food production – i.e. ensuring that sufficient food is being produced to feed the world's population ideally avoiding cost and damage to the natural environment. The core source of food is agriculture. A major challenge to sustainable food production is climate change. Climate change has caused considerable change to the world's environment, and this has had an impact on agriculture.

One can find publicly available agricultural datasets that provide crop statistics over a long period of time – e.g. Kaggle provides such a dataset originating from the Food and Agriculture Organisation of the United Nations (FAO) providing yearly agricultural output from all countries in the world between 1961 and 2018. Such datasets can also be provided locally by the Government's Department of Agriculture.

This project is to analyse and mine such agricultural output datasets, and aggregate together with publicly climate data to investigate the impact of climate change on different crops. Such analysis allows the use of AI techniques to predict crop outputs, and provide forecasts for different scenarios of climate change.

Recommended Resources:
https://www.kaggle.com/datasets/raghavramasamy/crop-statistics-fao-all-countries
https://ourworldindata.org/agricultural-production

Ethical Issues:
N/A – this dissertation is envisaged to make use mainly of publicly available datasets

Title: Reinforcement Learning for Partially Observable Stochastic Games (Reinforcement Learning)

Supervisor: Dr Bajada Josef

Description: In this project we will explore the use of RL algorithms to create intelligent agents that play partially observable stochastic games. Partial observability means that the agent does not have full information of the current environment state. This could be due to hidden opponent information (in case of multi-agent games), or concealed information that still needs to be discovered. Stochasticity means that there is an element of chance or randomness in the game, and actions are non-deterministic (albeit, there might be a known probability distribution that governs their outcomes). Such games present various challenges, such as state representation, action encoding, action masking and selecting the right reward signals to guide algorithms to learn an effective policy. Such games typically present huge state spaces and action spaces, which are not solvable using traditional search-based methods.

For this project, one game will be selected, and the student will be expected to build a simulator for the game, and build a framework through which a policy to play the game well will be learnt. The student will be expected to review existent literature on similar games, and take inspiration from their ideas how to model the right state information, encode the possible actions to take in each state, craft the rewards for each action, and integrate the environment with reinforcement learning algorithms. The performance of these algorithms will be analysed and compared, as part of the evaluation of the project.

Recommended Resources:
Sutton, R. S. and Barto, A. G., "Reinforcement Learning: An Introduction", 2nd Edition, The MIT Press, 2018.

Farama Gymnasium: https://gymnasium.farama.org/
Farama PettingZoo: https://pettingzoo.farama.org/
RLLib: https://www.ray.io/rllib

Ethical Issues:

Title: Intelligent Demand-Control Policies for Smart-Grid Power Networks (Reinforcement Learning)

Supervisor: Dr Bajada Josef

Description: Future power networks have two challenges ahead of them. The first is to respond to higher demand due to electrification of services, such as transportation. The second is the effective use of renewables, which are not directly controllable, and only predictable within a short timeframe (often of a few hours). With electricity storage facilities being largely inefficient (70% at best), an alternative solution is to control and shift flexible demand, such as dishwashers or charging of electric vehicles, to times of the day of low aggregate consumption, or times where high surplus of renewable generation is predicted. The effect of this would be to flatten demand peaks and reduce the need for conventional power generation, which is often powered by fossil fuels.

In this project, the student will first review existent power demand simulators, to generate household load profiles and simulate realistic electricity demand. State-of-the-art demand-control policies will be reviewed and used as a baseline to benchmark the current solutions. AI-based solutions will be investigated, such as using Deep Reinforcement Learning with Continuous Control actions, to learn policies that schedule flexible demand, with the objective of flattening peaks and creating a more stable load profile. Time-permitting, renewable generation profiles will also be integrated in the mix, to generate a policy which prefers the use of clean energy versus conventional polluting sources.

Recommended Resources:
Bajada, J.; Fox M.; and Long. D. "Load modelling and simulation of household electricity consumption for the evaluation of demand-side management strategies." IEEE PES ISGT Europe 2013. IEEE, 2013.

Proedrou, Elisavet. "A comprehensive review of residential electricity load profile models." IEEE Access 9 (2021): 12114-12133.

Velosa, N.; Gomes, E.; Morais, H.; and Pereira, L. "PROCSIM: An Open-Source Simulator to Generate Energy Community Power Demand and Generation Scenarios". Energies 2023, 16, 1611.

Sutton, R. S. and Barto, A. G., "Reinforcement Learning: An Introduction", 2nd Edition, The MIT Press, 2018.

Ethical Issues:

Title: Electric Vehicle Assignment and Routing for Carpooling Fleets (AI Planning, Reinforcement Learning, Discrete Optimisation)

Supervisor: Dr Bajada Josef

Description: Autonomous carpooling services promise to democratize personal transportation through reduced costs, increased safety from reduced car accidents, less traffic and less need for parking spaces. With a fully electrified fleet, pollution is also bound to be reduced. For such services to be successful, it is important that the right sequence of pickup and drop-off activities is chosen for each vehicle, such that customers do not incur long waiting times or long travel times, while ensuring that the electric vehicle does not end up without enough battery charge.

In this project, the student will develop a system which, given a number of pickup and drop-off waypoints from different customers, will assign the right sequence of actions for each car. There will also be a number of charging points available, and the car can be instructed to drive to a charging point, as long as there are no more customers on board and the battery level is below a certain threshold. A number of approaches will be investigated, including traditional discrete optimization techniques such as Tabu search and Genetic Algorithms, together with more recent experimental approaches, such as using Reinforcement Learning. The results from this thesis will be a comprehensive analysis of the performance of different algorithms as scenarios become more complex with more customer trips and more vehicles to schedule accordingly.

Recommended Resources:
Lin, Bo, Bissan Ghaddar, and Jatin Nathwani. "Deep reinforcement learning for the electric vehicle routing problem with time windows." IEEE Transactions on Intelligent Transportation Systems 23, no. 8 (2021): 11528-11538.
https://lin-bo.github.io/files/DR_EVRPTW.pdf

Raza, Syed Mohib, Mohammad Sajid, and Jagendra Singh. "Vehicle Routing Problem using Reinforcement Learning: Recent Advancements." Advanced Machine Intelligence and Signal Processing. Singapore: Springer Nature Singapore, 2022. 269-280.
https://www.researchgate.net/profile/Mohammad-Sajid-12/publication/361549456_Vehicle_Routing_Problem_Using_Reinforcement_Learning_Recent_Advancements/links/636293ef54eb5f547c993f02/Vehicle-Routing-Problem-Using-Reinforcement-Learning-Recent-Advancements.pdf

Open Source Routing Machine: https://project-osrm.org/
Sutton, R. S. and Barto, A. G., "Reinforcement Learning: An Introduction", 2nd Edition, The MIT Press, 2018.

Ethical Issues:

Title: Ethics in AI - A systemic review of the Literature (AI Ethics)

Supervisor: Dr Camilleri Vanessa

Description: The scope of this FYP is to illustrate a systemic review of the literature concerning AI Ethics in particular aspects involving LLMs, VR and AR. These include legal acts and official publications about expert recommendations. As part of the methodology, a group of experts in AI Ethics and Law, will be interviewed to gain deeper insights into the complexities of the ethical implications of AI. Following this and an in-depth discussion of the literature, it is expected that the outcome from the FYP would be a comprehensive list of recommendations, including addressing any gaps identified in the literature and publications reviewed.

Recommended Resources:
The AI Act / European Commission Council Recommendations

Ethical Issues:
N/A

Title: VR Comm - Communication in VR using LLM (VR, LLMs)

Supervisor: Dr Camilleri Vanessa

Description: This FYP will integrate the use of LLMs in a VR environment to create a rich communication medium between humans and agents. The agent avatar in the VR will take on human like communication mechanisms that include the use of language (both text and voice) as well as gestures, to maintain a flowing conversation in context. The outcome from the FYP will be human-like avatar in VR that will potentially be able to sustain a 2-way situated conversation with a human user. This will be followed by a set of guidelines for the setting up and use of the human-like conversing agent, as well as a framework for the ethical practices in setting up such an agent in VR.

Recommended Resources:
Unreal Engine / selected publications

Ethical Issues:
N/A

Title: Avatars, Embodiment and Cognition (Cognitive Science, AI)

Supervisor: Dr Camilleri Vanessa

Description: This FYP will explore and discuss in-depth issues related to full body cognition in VR and VWs. It is expected that as part of the FYP, students will design a contextual VW in a VR setting, where the user is assigned an avatar that may or may not represent his/her true self. PCG may be used as part of the VW simulation. The user will be monitored and bio sensing measures will be taken to understand the depth of the effect of the VR environment and the situation created on the user's following actions and behaviour. This FYP will be delving into aspects of neuroscience, and will explore issues related to ethics and ethical practices for VR-based simulations.

Recommended Resources:
Unreal Engine

Ethical Issues:
Ethical practices will be followed, users will sign up voluntarily and safety will be ensured for all individual/s participating.

Title: User Engagement in Serious Games (User Profiling, Decision Making)

Supervisor: Dr Camilleri Vanessa

Description: This FYP proposal aims to develop a proof of concept for a serious game that adapts to different learning practices. The project focuses on investigating user engagement in serious games and its impact on successful learning outcomes. The main research questions to be addressed are: -Which user models are most effective in enhancing the interactive learning experience within a serious game? -What variables should be included in a user model for optimal learning outcomes? -What methods can be employed to accurately model the user within the serious game? Considering the growing popularity of AI in games and the increasing utilization of serious games for training and professional development, this project holds significant relevance and importance. By addressing these research questions, the project aims to contribute to the development of more effective and engaging serious games for learning purposes.

Recommended Resources:

Ethical Issues:

Title: Brain-to-Text (Natural Language Processing)

Supervisor: Prof. Dingli Alexiei

Description: The primary aim of this project is to develop a simple, cost-effective brain-computer interaction (BCI) system using commercially available EEG bands and deep learning techniques to map brain signals into words, thereby enabling non-invasive communication. Brain-computer interaction has the potential to revolutionise how humans communicate with each other and interact with technology. Current research in BCI often requires complex, expensive equipment and highly specialised knowledge. This project explores the feasibility of developing an accessible, user-friendly BCI system by leveraging commercial EEG bands and deep learning models.

Literature Review: Conduct a thorough review of existing research on EEG-based BCIs and deep learning methods to identify the most promising techniques and challenges.

Data Collection: Use a commercial EEG band to collect brain signal data from multiple participants while performing various cognitive tasks, such as thinking of specific words or phrases.

Data Preprocessing: Clean and preprocess the collected data to remove noise and extract relevant features for further analysis.

Model Development: Train a deep learning model, such as a recurrent neural network (RNN), long short-term memory (LSTM) network or others, to map brain signals into words or phrases. Optimise the model using cross-validation and appropriate performance metrics.

Evaluation: Test the performance of the developed model to evaluate its accuracy and generalizability. Analyze the results and identify potential improvements or limitations.

This project aims to produce a proof-of-concept BCI system that can accurately map brain signals into words using a commercial EEG band and a deep learning model. The results will provide valuable insights into the feasibility of using affordable, accessible technology for BCI research and pave the way for further exploration of this exciting field.

Recommended Resources:

Online resources such as

https://cns.utexas.edu/news/podcast/brain-activity-decoder-can-reveal-stories-peoples-minds?fbclid=IwAR3rHTh8FYBt_Orb2Y-t2b1FVRF-7_oxTzrSN6bmrI7-QBzAzqhnYXbroDw

https://medium.com/svilenk/bciguide-246a9ca76fcd

https://towardsdatascience.com/a-beginners-guide-to-brain-computer-interface-and-convolutional-neural-networks-9f35bd4af948

https://medium.com/geekculture/understanding-brain-computer-interfaces-with-python-209f32e10db8

Ethical Issues:

Adults will be used for the study, and we'll collect consent from each one.

Title: Large Language Model for Maltese (Natural Language Processing)

Supervisor: Prof. Dingli Alexiei

Description: The primary goal of this project is to develop a prototype large language model (LLM) for the Maltese language, a low-resource language that presents unique challenges in natural language processing. By addressing these challenges, this project aims to enrich Maltese language resources and improve the accessibility of digital services and tools for Maltese speakers.

Collect and curate a diverse dataset of Maltese text from various sources, such as websites, social media, newspapers, and literary works.

Preprocess and clean the dataset to remove inconsistencies and improve data quality.

Develop a simple LLM using state-of-the-art NLP techniques and tools, such as Transformers and tokenisation methods adapted for Maltese.

Evaluate the LLM's performance using appropriate metrics.

Challenges:
Data scarcity: As a low-resource language, Maltese has limited digital textual resources. This scarcity will require the diligent collection and curation of diverse data sources.
Orthographic variation: Maltese's orthographic variability, coupled with limited standardization, may impede the LLM's performance and necessitate additional preprocessing steps.
Model customization: Adapting existing NLP techniques and tools to accommodate the unique linguistic characteristics of Maltese will be a central challenge.

Expected Outcomes:
Upon completion, the project will yield a simple Maltese LLM that can serve as a foundation for future NLP applications, such as machine translation, sentiment analysis, and text summarisation, catering to the Maltese-speaking community. The project will also provide valuable insights into the challenges of working with low-resource languages and contribute to the ongoing conversation in the field of AI and NLP.

Recommended Resources:
Available online:

https://aclanthology.org/2021.mrl-1.11/

Ethical Issues:
None

Title: Academic Thesis Evaluator (ATE) (Natural Language Processing)

Supervisor: Prof. Dingli Alexiei

Description: The goal of this project is to develop a simple, user-friendly AI system that qualitatively evaluates the quality of undergraduate academic theses, ensuring the presence of all required elements and consistency throughout the document. Moreover, the system will employ machine learning models to predict the final grade of the thesis, which will be validated using a dataset of theses available in the university library.

The proposed Academic Thesis Evaluator (ATE) system will be built using Natural Language Processing (NLP) techniques and machine learning algorithms. The system will be designed to perform the following tasks:
Check for the presence and consistency of all essential thesis elements, such as the title page, abstract, table of contents, introduction, methodology, results, discussion, conclusion, and bibliography.
Analyse the coherence and consistency of the thesis, including the logical flow of ideas, argumentation quality, and adherence to the chosen citation style.
Employ machine learning models to predict the final grade of the thesis based on the quality assessment.
The ATE system will be trained and evaluated using a dataset of undergraduate theses available in the university library. It will be used as ground truth for the machine learning models.

Expected Outcomes:

A functional AI-driven system capable of qualitatively evaluating undergraduate theses' quality and predicting their final grades.
A user-friendly interface allows students and advisors to upload their thesis and receive feedback on the quality assessment and grade prediction.
An evaluation report detailing the system's performance, including accuracy, precision, recall, and F1 scores for grade predictions.

A user guide to help students and advisors navigate the ATE system and understand its output.

This project will provide undergraduate students with a valuable tool for self-assessment and improvement of their academic thesis, fostering a higher standard of research and writing across disciplines.

Recommended Resources:
Various online.

Ethical Issues:
None.

Title: Developing a protocol for human motion capture using wearable inertial sensors (Biomechanics, IoT)

Supervisor: Dr Galea Ingrid

Description: This project shall review the state-of-the-art in motion capture sensors, with focus on inertial sensors. Subsequently, the student shall develop a protocol for motion capture from inertial sensors. The protocol shall include calibration of the sensors, on-body sensor positioning, and normalisation methods for different participants. The final deliverable of this project shall be a sample dataset acquired using the proposed protocol.

Recommended Resources:
https://www.xsens.com/hubfs/Downloads/Software/DOT/Educational%20curriculum/Xsens_DOT_Educational_curriculum.pdf?hsCtaTracking=3974334a-5a1e-4cc5-8753-fef65f73c81d%7Cd669e062-a9b1-46f5-99fa-5c79c38d4ad0

Ethical Issues:
Ethical consent is to be sought when acquiring motion data from human participants.

Title: Developing a biomechanical model for motion analysis of the human lower limb (Biomechanics, IoT)

Supervisor: Dr Galea Ingrid

Description: This project shall review the main biomechanical models that are used to represent human motion of the lower limb. Such models are typically used in human gait analysis. The project shall subsequently compute the joint angle given accelaration data acquired from inertial sensors positioned on the limbs. This project shall use readily available datasets or simulated data.

Recommended Resources:
https://www.matec-conferences.org/articles/matecconf/pdf/2022/04/matecconf_mms2020_03006.pdf

Ethical Issues:

Title: Visualisation of inertial data from wearable sensors (Biomechanics, IoT)

Supervisor: Dr Galea Ingrid

Description: This project shall review how motion capture data can be visualised. It shall then create animations that represent motion data in three different reference frames: (1) a local reference frame (2) a segment reference frame and (3) a joint reference frame. Visualisations are to be customisable according to varying sensor positioning and body measurements. This project shall use readily available datasets or simulated data.

Recommended Resources:
https://base.xsens.com/s/article/MVN-Biomechanical-Model?language=en_US

Ethical Issues:

Title: An Investigation of Context-Free Language Learning Techniques (Machine Learning, Grammatical Inference, Search Algorithms, Heuristics, Algorithm Design)

Supervisor: Dr Guillaumier Kristian

Description: Grammatical inference is the task of learning a formal grammar from strings which belong to a language and strings which do not belong to a language. This inference task has many real-world applications including robotics, data mining, structural pattern recognition, speech recognition, and bioinformatics.

In this project, we are concerned with surveying and implementing techniques used in the inference of context-free languages (CFL) from structured data. The primary aims are to:
- Survey the literature regarding current state-of-the-art algorithms.
- Investigate practical applications of CFL inference.
- Build a simple framework consisting of foundational data structures and CFL learning algorithms.
- Verify the behaviour of our implementations on Omphalos-style problems.

Recommended Resources:
Material and tutorials will be provided to the student to support the FYP

Ethical Issues:
None

Title: DFA Learning using Neural Networks (Machine Learning, Grammatical Inference, Search Algorithms, Heuristics, Algorithm Design)

Supervisor: Dr Guillaumier Kristian

Description: Grammatical inference is the task of learning a formal grammar from strings that belong to the language and strings that do not. In this project, we will focus on the identification of regular languages (as Deterministic Finite State Automata, DFAs) from training sets consisting of both positive and negative examples. This inference task has many real-world applications including robotics, data mining, structural pattern recognition, speech recognition, and bioinformatics.

While state merging algorithms are one of the most powerful and successful approaches in DFA learning, there has been increasing interest in exploring connectionist approaches to the problem (typically RNNs, LSTMs, and transformers). In this FYP, we aim to study how neural networks can be used to learn DFAs (both in terms of training training and DFA extraction), what their advantages and limitations are, and compare their behaviour to state merging approaches.

Recommended Resources:
Material and tutorials will be provided

Ethical Issues:
None

Title: Extending DfaGo with Python (Machine Learning, Grammatical Inference, Search Algorithms, Heuristics, Algorithm Design)

Supervisor: Dr Guillaumier Kristian

Description: Grammatical inference is the task of learning a formal grammar from strings that belong to the language and strings that do not. In this project, we will focus on the identification of regular languages (as Deterministic Finite State Automata, DFAs) from training sets consisting of both positive and negative examples. This inference task has many real-world applications including robotics, data mining, structural pattern recognition, speech recognition, and bioinformatics.

One of the most important prerequisites before one can start working on any DFA learning project, is to develop all the underlying data structures, search algorithms, and analysis tools efficiently. A lot of this work has already been done in an existing DFA learning framework called DfaGo. The aim of this FYP is to extend DfaGo using Python. The idea is to create an easy-to-use, low-overhead, and scaleable interface that integrates Python's powerful machine learning frameworks with DfaGo. This will allow for the implementation of algorithms based on 'non-classical' DFA learning paradigms such as neural networks and transformers.

Recommended Resources:
Material and tutorials will be provided

Ethical Issues:
None

Title: Learning to rank humans' emotional states (Affective computing, Machine learning)

Supervisor: Dr Makantasis Konstantinos

Description: Recognising humans' emotional states is a challenging task and arguably one of the most fundamental in affective computing. From a computational point of view, the main challenge stems from the fact that emotions are subjectively defined. Psychological theories and evidence from multiple disciplines suggest that treating emotions as ordinal variables better aligns with their subjective nature. Standard emotion recognition practice, however, considers subjectivity a nuisance and tries to engineer it away. In other words, it treats emotions as nominal variables and approaches the problem of emotion recognition through a classification or regression perspective. This study aims to explore different ranking algorithms for recognising humans' emotional states. Ranking algorithms will treat emotions as ordinal variables and, thus, respect their subjective nature. Towards this direction, several ranking algorithms, such as ordinal logistic regression, ordinal neural networks, RankNet and LambdaRank algorithms, will be implemented and applied to the problem of emotion recognition. The implemented algorithms will be evaluated and compared against state-of-the-art in terms of computational efficiency, sample complexity, and recognition accuracy using two publicly available datasets (AGAIN and RECOLA or SEWA). The outputs of this study will be i) a thorough evaluation of ranking algorithms in the framework of emotion recognition and ii) an open-access code repository for downloading the implemented algorithms.

Recommended Resources:
[1] Yannakakis, Georgios N., Roddy Cowie, and Carlos Busso. "The ordinal nature of emotions: An emerging approach." IEEE Transactions on Affective Computing 12.1 (2018): 16-35.
[2] Burges, Christopher JC. "From ranknet to lambdarank to lambdamart: An overview." Learning 11.23-581 (2010): 81.
[3] Melhart, David, Antonios Liapis, and Georgios N. Yannakakis. "The arousal video game annotation (AGAIN) dataset." IEEE Transactions on Affective Computing 13.4 (2022): 2171-2184.

Ethical Issues:
N/A

Title: Semi-supervised Learning for Affect Modelling (Affective computing, Machine learning)

Supervisor: Dr Makantasis Konstantinos

Description: Affective computing strives to unveil the unknown relationship between affect elicitation, the manifestation of affect, and affect annotations. The ground truth of affect, however, is predominately attributed to the affect labels, which inadvertently include biases inherent to the subjective nature of emotion and its labelling. The response to such limitations is usually augmenting the dataset with more annotations per data point; however, this is rarely possible. This study aims to reframe affect modelling by investigating the degree to which knowledge from unlabelled data can boost the performance of affect modelling. Towards this direction, this study will leverage tools from the semi-supervised learning paradigm. Semi-supervised concepts, such as co-training and graph-based label propagation, will be implemented to inject knowledge about the distribution of unlabelled data points into the models of affect. The implemented models will be evaluated against fully supervised learning models in terms of prediction accuracy, sample complexity and computational efficiency using two publicly available affect corpora (AGAIN and RECOLA or SEWA). The outputs of this study will be i) a thorough evaluation of semi-supervised algorithms in the framework of affect modelling and ii) an open-access code repository for downloading the implemented methodologies. Successfully applying semi-supervised methods to affect modelling will be a significant result for affective computing and emotional artificial intelligence, making a step towards the application of affect models in-the-wild.

Recommended Resources:
[1] Yang, Xiangli, et al. "A survey on deep semi-supervised learning." IEEE Transactions on Knowledge and Data Engineering (2022).
[2] Zhang, Jianhua, et al. "Emotion recognition using multi-modal data and machine learning techniques: A tutorial and review." Information Fusion 59 (2020): 103-126
[3] Melhart, David, Antonios Liapis, and Georgios N. Yannakakis. "The arousal video game annotation (AGAIN) dataset." IEEE Transactions on Affective Computing 13.4 (2022): 2171-2184.

Ethical Issues:
N/A

Title: From Context-specific to Context-agnostic Player Affect Modelling (Affective computing, Machine learning)

Supervisor: Dr Makantasis Konstantinos

Description: Despite the advances in Affective Computing, avowedly based on today's powerful deep-learning algorithms, the scope of the derived affect models is still narrow. The current state-of-the-art suggests that any success of affect modelling is heavily dependent on the user, the task at hand, and the context in general. This specificity limitation is detrimental both for the scientific value and for the practical applicability of the methods developed and studied in the affective computing field. As affect modelling has predominately been tested on single tasks under a specific context, there is little to be said about the generality of the models obtained. This study aims to overcome the specificity limitation by building on the concepts of supermasks and dynamic convolutions. Specifically, supermasks and dynamic convolutions will allow, first, the development of accurate context-specific models and, then, their fusion into a unified context-agnostic model. This study will investigate the degree to which context-specific knowledge can be transformed into context-agnostic knowledge for cross-task and cross-user applicability. Given that digital games offer complex yet well-defined and controlled problems for exploring the efficiency of general AI, they define the ideal arena for testing the generalization capacity of the developed models. The outputs of this study will be i) a thorough investigation of the application of supermasks and dynamic convolutions for building context-agnostic models of affect and ii) an open-access code repository for downloading the implemented methodologies. Successfully transferring context-specific knowledge to context-agnostic models will be a significant result for affective computing and emotional artificial intelligence, making a step towards the application of affect models in-the-wild.

Recommended Resources:
[1] Makantasis, Konstantinos, Antonios Liapis, and Georgios N. Yannakakis. "The pixels and sounds of emotion: General-purpose representations of arousal in games." IEEE Transactions on Affective Computing (2021).
[2] Melhart, David, Antonios Liapis, and Georgios N. Yannakakis. "The arousal video game annotation (AGAIN) dataset." IEEE Transactions on Affective Computing 13.4 (2022): 2171-2184.
[3] Wortsman, Mitchell, et al. "Supermasks in superposition." Advances in Neural Information Processing Systems 33 (2020): 15173-15184.
[4] Chen, Yinpeng, et al. "Dynamic convolution: Attention over convolution kernels." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2020.

Ethical Issues:
N/A

Title: Continual Learning for Hyperspectral Data Classification (Remote Sensing, Machine learning)

Supervisor: Dr Makantasis Konstantinos

Description: The spectrum in many narrow spectral bands through hyperspectral imaging provides valuable information towards material and object recognition. Current state-of-the-art models approach the problem of material and object recognition as a pixel-wise classification problem and try to build highly accurate classifiers for specific datasets. This implies that the developed models should be retrained (usually from scratch) when the sensing device changes (e.g., satellites with different hyperspectral sensors) or the classes of materials are enhanced (e.g., the model learns to classify X different materials, but the new data contain one more material). This study will leverage tools from the continual learning paradigm to develop "extensible" hyperspectral data classifiers. Specifically, the classifiers will be trained on a dataset captured with a specific sensor and depicting a known number of materials. Then, based on continual learning methodologies, the classifiers will be able to adapt to new datasets captured with different sensors and depicting different materials. The developed methodologies will be tested against state-of-the-art models in terms of prediction accuracy, sample complexity and computational efficiency using publicly available remote sensing corpora (Indian Pines, Salinas, Pavia Centre, Pavia University, Kennedy Space Center and Botswana). The outputs of this study will be i) a thorough evaluation of continual learning algorithms for pixel-wise hyperspectral data classification and ii) an open-access code repository for downloading the implemented methodologies. Building classifiers able to adapt to new datasets will be a significant step towards the derivation of general (universal) models for hyperspectral data classification.

Recommended Resources:
[1] Makantasis, Konstantinos, et al. "Deep supervised learning for hyperspectral data classification through convolutional neural networks." 2015 IEEE international geoscience and remote sensing symposium (IGARSS). IEEE, 2015.
[2] De Lange, Matthias, et al. "A continual learning survey: Defying forgetting in classification tasks." IEEE transactions on pattern analysis and machine intelligence 44.7 (2021): 3366-3385.
[3] Hadsell, Raia, et al. "Embracing change: Continual learning in deep neural networks." Trends in cognitive sciences 24.12 (2020): 1028-1040.

Ethical Issues:
N/A

Title: VWLE - Virtual World Learning Environment (AI in Education)

Supervisor: Prof Montebello Matthew

Description: The future of VLE is the use of AI to optimise the learning process. VWs are the next medium to employ for students to take full advantage of available educational resources while injecting smart AI processes to add value.

Recommended Resources:
https://www.igi-global.com/chapter/vle-meets-vw/273045
https://axonpark.com/
https://www.spatial.io/

Ethical Issues:

Title: iTrack - An intelligent running track for athletic analytics (IoT)

Supervisor: Prof Montebello Matthew

Description: Envisaging the next generation smart running track involved the use of multiple sensors but most of all of a purposely equipped track supported by a robust AI backend that provides personalised information.

Recommended Resources:
https://smartracks.run/
https://www.polytan.com/smart/running-tracks/

Ethical Issues:

Title: Deep Marine Archaeology (Autonomous Underwater Drones)

Supervisor: Prof Montebello Matthew

Description: Collecting archaeological data underwater at depths too dangerous for humans to access is an ideal setting for smart autonomous vehicles. This project incorporates the use of specialised hardware together with intelligent collection of data and real-time decision making.

Recommended Resources:
http://web.mit.edu/deeparch/www/publications/papers/MindellCroff2002.pdf
https://www.academia.edu/6116313/OTC_18841_Deepwater_Archaeology_With_Autonomous_ Underwater_Vehicle_Technology

Ethical Issues:

Title: Multi Drone Intelligence (Multi-agent systems)

Supervisor: Prof Montebello Matthew

Description: The applications involving multiple software agents communicating together to efficient get a job done vary from military use to entertainment. This project investigates the use of such setups and implements one as a proof of concept.

Recommended Resources:
https://multidrone.eu/2020/01/23/multidrone-feature/
https://www.researchgate.net/publication/337260183_Multiple_UAV_Systems_A_Survey

Ethical Issues:

Title: HELP - Health Ethical Learning Program (Ethics & AI)

Supervisor: Prof Montebello Matthew

Description: The main objective of this project is to develop a machine learning tool based on artificial intelligence that supports health care professionals in identifying, verbalising, articulating and making judgements about ethics in clinical practice.

Recommended Resources:
Lukas J. Meier, Alice Hein, Klaus Diepold & Alena Buyx (2022). Algorithms for Ethical Decision-Making in the Clinic: A Proof of Concept, The American Journal of Bioethics, 22:7, 4-20, DOI: 10.1080/15265161.2022.2040647

Danton Char (2022) Important Design Questions for Algorithmic Ethics Consultation, The American Journal of Bioethics, 22:7, 38-40, DOI: 10.1080/15265161.2022.2075054

Lukas J. Meier, Alice Hein, Klaus Diepold & Alena Buyx (2022) Clinical Ethics – To Compute, or Not to Compute?, The American Journal of Bioethics, 22:12, W1-W4, DOI: 10.1080/15265161.2022.2127970

Ethical Issues:

Title: Optimisation of Saliency-driven Image Content Ranking Parameters (Computer Vision)

Supervisor: Dr Seychell Dylan

Description: Saliency-driven image content ranking (SaRa) is an approach that automatically ranks the selection of objects based on their visual saliency in the image context without needing to pre-train a model. This modern approach in computer vision allows for prioritising visual content before further processing. SaRa allows regions in an image to be selected according to the desired rank based on the analysis of different regions of the image. The image is split into a grid, and each segment is processed. In the original publication, an experiment was conducted to identify an optimal fixed size for the grid. However, different grid sizes can potentially have different results on content ranking. The resolution of the image might also have an impact on these results. Based on these hypotheses, this project is an experiment that looks into the optimisation of parameters in this technique and a subsequent investigation into their impact on the output result.

Recommended Resources:
D. Dr Seychell and C. J. Debono, "Ranking Regions of Visual Saliency in RGB-D Content," 2018 International Conference on 3D Immersion (IC3D), Brussels, Belgium, 2018

Ethical Issues:
NA

Title: Decomposition and Analysis of Inset Images in news articles. (Computer Vision)

Supervisor: Dr Seychell Dylan

Description: The analysis of images in news articles is important in the investigation of potential bias in the article. The choice of image and the headline's text and caption can easily mislead the reader and induce bias. This automatic investigation of such bias, formally called picture selection bias, is rendered trickier when the primary image in the article contains other inset images. For example, an article about an enterprise can contain an image of its headquarters with an inset image of a person. The choice of which images are inset is important because it can easily change the context or bias of the article. In the above example, the person could be the CEO, a politician, or someone vaguely associated with the enterprise. Each of these alternatives would lead the reader to different perceptions. This project aims to classify and subsequently decompose inset-images into a subset of images that can then be analysed in relation to the caption or the related article. While the process can be fully automated, user studies can also be included in the process of this project.

Recommended Resources:
All necessary resources will be provided by the supervisor.

Ethical Issues:
NA

Title: Investigation of Visual Bias in Generative AI (Computer Vision)

Supervisor: Dr Seychell Dylan

Description: The outstanding results of Generative AI in creating new visual content, primarily images, have attracted the attention of various individuals. Platforms such as Mid-Journey, DALL-E2 and Stable Diffusion have empowered many individuals to create custom content based on text prompts and use the output in various contexts. This ranges from entertainment and arts to the generation of images for articles. The output could be biased, depending on the data in the training set and/or in the latent space that the diffusion model uses in the generation phase. This project aims to investigate the bias of Generative AI by focusing on the open source and queryable dataset LAION-5B that is used to train the StableDiffusion model. Given a prompt and the stochastic nature of the model, different output results can be investigated for bias. The main elements of the prompt will also be used to investigate the images in the dataset. This will provide insight into whether there is any bias in the output and, if so, from where it originated. For example, the prompt "A doctor having a conversation with a nurse" can be used. The project will use existing and evaluated computer vision techniques to classify the gender of the doctor and nurse at the output. The same can be done on the images in the dataset for the key entities, in this case, "doctor" and "nurse". Since the prompt did not specify the gender of either entity, the output should be fairly balanced in the choice of gender.

Recommended Resources:
Schuhmann et al. "LAION-5B: An open large-scale dataset for training next generation image-text models." NeurIPS (2022)

Ethical Issues:
NA