School of Computer Science

Research Projects

PhD Scholarships 2024

Updated on 15 Jan 2024, new projects added

Artificial Intelligence and Machine Learning

Bioinformatics and Health Informatics

Emerging Topics

Foundation of Computing

Human Computer Interaction and Multimedia

Security and Networks

Software Engineering and Distributed Systems

ADDED PROJECTS – JAN 2024

Artificial Intelligence and Machine Learning

01. Examining the Interplay between Time Series Classification and Explanation

This project aims to explore the relationship between time series data and explainable AI. We will investigate how XAI can enhance the interaction between time series classifiers and model explanations, with the goal of improving both the model and dataset. Our approach involves creating a feedback loop that leverages XAI techniques to continuously refine the dataset and classifier. We seek to address the following questions: - How can we use explanation methods to enhance classification methods in terms of accuracy, efficiency and robustness? - How can we leverage explanation methods to improve or reduce the input data? - What is the optimal design for a feedback loop that optimises both data and classifier? - Which combination of classifier models and explanation methods yields the best performance? A wide variety of methods were developed to address time series classification (TSC). These methods are different in the way that they use the data to extract useful features. We have access to many datasets for TSC (UCR/UEA popular benchmarks used in the community), and there is clarity in the field of what types of classifiers are most accurate for different domains [1]. When it comes to XAI for time series data, this area is much younger and under active development. XAI was mainly used for image classification and natural language processing, with some recent efforts focusing on time series data. The focus so far has been to identify features or samples that make a large contribution to classifying the data. Recent work [2,3] focuses on evaluating multiple explanation methods for TSC, to identify the most useful explanations for a given dataset and classifier. However, that work does not use methods or metrics to create a better classifier and/or data. Based on this paper and others [4,5] we note that little work is done towards optimising the interaction between data, classifier and explanation.

A few papers proposed to research the interaction between XAI methods and data or classifiers [6,7]. However, these works are not related to time series classification.

In this project we focus on the relatively unexplored topic of leveraging XAI methods for data and model optimisation. Specifically, the novelty of this project is in the use of XAI to understand important aspects of data and model learning. This understanding can help us refine the data and model iteratively, ultimately leading to better data quality and model performance. For instance, in a neural network, changes could include altering the architecture, the feature extractor, or selecting different hyper-parameters. For creating better data, we can remove or discount noisy data, or extract important features.

[1]http://arxiv.org/abs/2304.13029

[2]https://doi.org/10.1038/s42256-023-00620-w

[3]https://arxiv.org/abs/2306.05501

[4]https://doi.org/10.1016/j.inffus.2023.101805

[5]https://doi.org/10.1007/s10618-023-00933-9

[6]https://doi.org/10.1007/s10618-023-00933-9

[7]10.1109/CVPRW59228.2023.00396

Keywords

Machine Learning, Time Series Classification, Explanation, Optimisation, Trustworthy AI

02. Give GANs to mobile devices for Split Learning enabled Federated Learning

Mobile devices possess multiple sensors that enable them to generate and collect datasets of great variety that are valuable when producing AI models because they increase diversity and accurately represent human activity. However, the owners of such devices are disincentivized to share their data due to privacy concerns. On the other hand, mobile users are inclined to participate in collaborative learning protocols that produce AI models to enhance the quality of their experience while using the produced models. The most popular collaborative learning protocol is federated learning (FL). Unfortunately, most devices currently in use cannot participate in FL protocols that train models of millions of parameters~(e.g., ResNet) because they do not have the computational and memory resources needed. Furthermore, in scenarios where the distribution of data between the devices varies significantly, the trained models have low accuracy. Techniques based on split learning (SL) reduce the resource needs of mobile devices by offloading the computationally heavy parts of the models to devices with higher capabilities.

The primary goal of this project is to develop an open-source framework that uses SL and GANs on mobile and cloud devices to train AI models of millions of parameters while, in parallel, solving the non-IID problem in FL. The framework will have a core component that will offer the basic functionality of SL on GANs to produce pre-trained models that generate IID synthetic datasets in a privacy-aware way and support the proposed protocol. The framework will be composed of two parts, the cloud and the mobile, and will be implemented upon existing well-known open-source FL frameworks. Some indicative examples are Flower, PySyft, and FATE.

The successful candidate will join a growing team of four PhD students, one postdoctoral fellow and one research assistant with research foci that overlap highly with this project. Additionally, this project aligns with one ongoing European project that applies AI to orchestrate cloud resources and one industrial project that applies AI to manage the resources of mobile devices. Additional opportunities include visiting international collaborators of the group in Ireland, the United Kingdom, Denmark, Greece, Hong Kong, and the United States of America.

Keywords

Generative AI, Split Learning, Mobile Computing, Privacy

03. Automating AI Threat Intelligence

AI is transforming the way we live and work. AI systems will soon seamlessly share and process information, aid businesses, governments and healthcare. This rapid integration of AI comes with challenges. Sometimes, AI systems will fail and generate potentially dangerous incidents, despite the risk management and AI governance controls envisioned by the EU AI Act. Similar to cybersecurity, where incidents cannot be eliminated, AI-related incidents will persist. To respond promptly, cyber threat intelligence (CTI) systems automate the sharing of incident data. Managing AI incidents would benefit from similar incident-sharing systems, and enable integrating AI incident handling with cyber incident handling. The OECD published an initial report seeking to define AI incidents in October 2023. Current approaches to AI incident tracking like AIAAIC (AI, Algorithmic, and Automation Incidents and Controversies) are based on human curation and sharing of documents. This limits machine-based processing and automated response to AI.

The primary aim of this research is to enhance the efficiency, accuracy, effectiveness of AI incident data sharing through by adapting existing cyber threat intelligence (CTI), sharing data formats and protocols, e.g. the open source MISP (Malware Information Sharing) platform, and the development of new metrics, tools and methods for AI incident classification, enrichment and normalisation.

The research questions guiding this project are:

Q1. To what extent can the efficiency of AI incident data sharing be improved by adapting standard cybersecurity CTI data formats and protocols and developing new AI incidents metrics or analytics?

Research Objectives

1. Define data formats, methods for processing AI incidents from distributed sources, ensuring compatibility with existing CTI standards and platforms.
2. Develop novel metrics to improve the speed and accuracy of AI incident data sharing while maintaining precision.
3. Evaluate new machine-readable AI Incident data sharing systems in lab conditions and a real-world context.
4. Engage with CTI standardisation, information and compliance bodies such as OASIS, ENISA, to promote and validate the research outputs.

Q2. *To what extent can AI incident data be accurately and effectively shared using adapted data models and tools, and how does this integration impact existing workflows and detection rates?*

Research Objectives

1. Adapt data models and tools within a CTI platform to incorporate diverse incident scenarios
2. Evaluate the ease of integration of AI incidents into existing MISP, CTI sharing, analysis workflows.
3. Measure the impact on detection rates for AI incident scenarios after the integration, indicating the effectiveness of the adapted platform.

The proposal spans AI incident management, cybersecurity, and data-sharing. A gap exists in machine-readable reporting in AI incidents, and AI-generated incidents are often overlooked in cybersecurity. Recent survey papers (2021) highlight the importance of cyber response due to CTI automation. Competitive standards like MISP data model and Mitre STIX necessitate unified approaches. This research aims to fill these gaps, paving the way for new advances in effectively governing AI to make it more trustworthy by improving its governance and enabling a more controlled and automated response to AI incidents.

Keywords

Trustworthy AI, AI Incident, Knowedge Extraction, Cybersecurity, Data Analytics

04. Adversarial AI approach for the verification and auditing of AI-generated content in education

"AI-generated content has been applied in many machine learning tasks such as multimodal translation, generation, summarising, captioning, data-to-any-modal generation, question-answer generation, and open-ended or long-form essay/story/video/audio creation. Commercial-grade models are readily available to everyone and accessible via API, such as ChatGPT, which can generate readable content with one-line prompts. In the context of academia, these technologies pose a significant threat as AI-generated content is increasingly becoming highly sophisticated, and it is getting harder to correctly identify, distinguish, and attribute the content to humans as it is easy to generate text that can trick a human into thinking another human wrote it.

Mainly in the assessment domain, establishing author attribution of content submitted by students for assessment in their documents is a new problem educators need to be equipped to address. Most of the anti-cheating and anti-plagiarism software available for universities compare student papers against a corpus of content to identify similarities. These techniques fail against AI-generated content as the text is not copied or similar from the corpus it is trained and is generated from a similar latent space where one has to establish attribution to one human author vs. several AI text generation models.

Detecting Large Language Model (LLM)-generated text is an important task in natural language processing (NLP) that helps distinguish between texts created by language models and those written by humans. This distinction ensures system security, prevents misuse, and fosters trust in new tech products. Various detection approaches have been proposed ranging from rule-based systems to machine learning-based models using syntactic, semantic, and pragmatic information. Techniques such as authorship attribution, style analysis, and anomaly detection are also common.

For educators, providing accurately personalised feedback to each student is challenging, especially in diverse assessment settings like essays, multiple-choice, or dynamic projects. Therefore, developing a system for automated detection and evaluation in a unified system presents significant technical challenges.

This research proposal aims to develop a new and effective adversarial AI-based approach to detect synthetic textual content generated by AI models. Centred research questions are as follows:

(i) What are the adversarial properties of text in the educational domain?

(ii) How to define metrics to evaluate whether the given text is human vs. machine-generated?

(iii) What is an efficient and robust black box adversarial method that can be applied to the passage being assessed, derived from the adversarial properties of text?(

(iv) How to explain the text generated by the AI model?"

Keywords

Large Language Model, Adversarial, Explainable AI

05. Machine Learning for Time Series Analysis of Human Running Data Using Accelerometers

Understanding the dynamics of walking and running has important implications for many health conditions where people struggle with walking as a result of illness and also for sport where the focus is on improving performance and reducing the risk of injury.

Gait analysis has been performed mostly using computer vision techniques. We will study the problem using inertial motion sensors which are placed on the body at strategic joints. The inertial motion sensors provide us with 3 dimensional accelerometer data with very high time resolution. State of the art machine learning and deep learning techniques need to be developed to understand these multi-dimensional time series and relate them to the mechanics of motion of the human body while walking and running.

The data for this project will be collected in association with Prof Kieran Moran's group in DCU. We have an initial data comprising accelerometer data, heart rate and motion capture data from 100 runners. We will continue this trial and recruit further participants as necessary.

Our goal is to develop methods which can analyse the accelerometer time series and extract features which help us predict potential issues in gait which could lead to injury or poor performance. Recent work in the field of time series analysis has lead to some new approaches for identifying and extracting patterns in multi-dimensional time series data. The development of efficient motif discovery algorithms for multi-dimensional time series of accelerometer data would be useful as a tool for summarising and visualising the important features which determine human gait and performance.

The analysis of accelerometer data allows to identify the important features for gait and body dynamics while walking and running. We intend to leverage these features to recommend new training methodologies, ways of preventing injury and personalised fitness programs for runners. This work will have impact on healthcare, sports science, biomechanics and the running community.

Keywords

Motif discovery, machine learning, injury prevention, gait analysis, time series modelling

06. Exploring Synergies between Quantum Computing and Machine Learning

Understanding how and when to leverage machine learning for and/or with quantum computing is still in its infancy. Machine learning can be used to optimize quantum algorithms, and quantum computing can be used to develop new machine learning algorithms. This topic will explore one of these directions (depending on the interest of the candidate). In quantum machine learning, i.e. using quantum computer to develop new machine learning methods, there are many challenges to solve: how to prepare data, developing new methods, interpreting the output of the quantum computer, handling error/noise etc. Similarly, when applying machine learning methods to quantum problems (e.g. controlling quantum systems, quantum error correction etc.) the challenge is to balance a rich understanding of the quantum system(s) with sufficient machine learning expertise to develop toolkits of value to the domain. Thus, the purpose of this PhD topic is to generate valuable tools and methods to further the understanding of where machine learning and quantum computing can benefit each other and identify meaningful application areas.

Keywords

Quantum Computing, Machine Learning

07. Enhancing Legal Text Understanding and Analysis: Exploring the capabilities of Natural Language Processing in the Legal Domain

Transformer-based Language Models (including Large Language Models) have revolutionised Natural Language Processing in recent years. The release of BERT [1] in 2018 and ChatGPT (https://chat.openai.com/) in 2022 in particular, have contributed to enormous growth in the capabilities of NLP technologies. Despite this, legal documents (e.g. legislation, treaties, court judgments, contracts, etc.) have their own particular characteristics that pose problems for modern approaches. These challenges include such issues as extremely long sequence lengths, specialised vocabulary, data imbalance [2]. Even modern transformer-based methods and Large Language Models (LLMs) still struggle with the characteristics of datasets such as EurLex (https://huggingface.co/datasets/eurlex) and LexGLUE (https://huggingface.co/datasets/lex_glue) [3].

This PhD project aims to address the specific difficulties in analysing legal texts, leveraging and advancing the capabilities of modern NLP techniques. The research will focus on identifying and mitigating the limitations of existing models to enable more effective applications within the legal domain.

Objectives:

1. Legal NLP State of the Art: - Investigate the applicability and limitations of current state-of-the-art methods within the legal domain.
2. Model Adaptation, Architecture Development, and Fine-tuning: - Assess the performance of existing Transformer models/LLMs on tasks relating to legal text. - Propose and implement adaptation techniques, (possibly including domain-specific fine-tuning, transfer learning, zero/few shot learning, etc.), to address shortcomings in current models and/or architectures.
3. Context-aware Analysis: - Explore methods to incorporate contextual information into legal text analysis. - Incorporate contextual information into the developed model/architecture to improve analytical performance.
4. Explainability and Interpretability: - Investigate methods to enhance the explainability and interpretability of model predictions in legal contexts (given that transparency is a key requirement of legal analysis).
5. Ethical and Legal Implications: - All work must be conducted in accordance to the ethical considerations of deploying advanced NLP models in the legal domain, particularly concerning biases, fairness, and accountability. Methodology: The research will involve a combination of literature reviews, dataset curation, model training and evaluation, and iterative development. The project will employ a multidisciplinary approach, drawing on expertise from both computer science and law.

References:

[1] Kenton, J.D.M.W.C. and Toutanova, L.K., 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of NAACL-HLT (pp. 4171-4186).

[2] Jayakumar, T., Farooqui, F., and Farooqui L. Large Language Models are legal but they are not: Making the case for a powerful LegalLLM, 2023, arXiv:2311.08890

[3] Mamakas, D., Tsotsi, P., Androutsopoulos, I., and Chalkidis, I. Processing Long Legal Documents with Pre-trained Transformers: Modding LegalBERT and Longformer. In Proceedings of the Natural Legal Language Processing Workshop 2022, pages 130–142, 2022, doi:10.18653/v1/2022.nllp-1.11

Keywords

Natural Language Processing, Large Language Models, Law

08. Computer Vision Based Indoor Multimedia Geolocation

The task of multimedia geolocation is becoming an increasingly essential component to effectively combat human trafficking and other illegal acts. While text-based metadata can easily provide geolocation information with access to the original media, this metadata is stripped when shared via social media and common chat application. Geolocating, geotagging, or finding geographical clues in the multimedia content itself is an arduous, complex tax. While there are numerous manual/crowdsourcing approaches to this, recent research has shown that computer vision is one viable avenue for research. The research on this project focuses on the curation of datasets for multimedia geolocation and developing novel computer vision-based techniques for indoor multimedia geolocation. The aim is to develop powerful methods for image geolocation that enable more efficient investigations in the field of human trafficking. As one example, colour values serve here as a key component to describe specific characteristics of an image and colour-based descriptors will be used for Content-Based Image Retrieval. The performance of the developed methods will be evaluated using the Hotels-50K dataset as an initial avenue of exploration before expanding to other diverse indoor multimedia datasets.

Keywords

Computer Vision, Artificial Intelligence, Digital Forensics

09. News Recommender Systems: Going Beyond Accuracy

News recommender systems (NRS) have been widely applied for online news websites to help users find interesting articles. The majority of existing work in this area focuses on predicting relevance between users and news articles e.g. by optimizing Click-Through Rate (CTR), but not paying enough attention to the quality of news articles recommended. Today‘s news platforms collect news articles from multiple sources, e.g., Google News watches more than 50k news sources worldwide. The quantity of news articles is huge, but the quality of these articles varies. Recommending Low-quality news articles, e.g., clickbait articles (39.26 % in unreliable media 1) that have sensational titles but no substantive information, may damage new platforms’ trustworthiness, and favorability. This project aims to build novel solutions to recommend users not only relevant but also “high-quality” news articles to improve their long-term satisfaction rather than only focusing on their short-term CTR. Many different solutions will be explored to identify “high quality” news articles: e.g., labelling news articles as “high quality” or “low quality” by taking into account the average dwell time. It is assumed that articles clicked frequently but with short dwell time are of lower quality. These labelled data will then be used to train a quality classifier that can be used to identify the quality of novel articles without click and dwell time information or filter news articles. Design complex cost functions taking into account CTR and quality and balancing their trade-off. In addition to the quality, the project will explore other factors, such as newsworthiness, diversity (recommended items are not similar) etc., and design a framework to evaluate our solutions properly.

Keywords

Recommender Systems, News Articles, NLP

10. Generative models for the development of RNA based therapeutics

Generative models are transforming many fields, including biology (Wainberg et al, Nature Biotechnology 2018). Recently generative models have been used to design new functional molecules including peptides, proteins and antibodies (Das et al, Nature Biomedical Engineering 2021; Anishchenko et al, Nature 2021; Shuai et al, Cell Systems 2023). Models for RNA have received less attention, although RNA based therapeutics hold significant promise (Zhu et al, Cell Death and Disease, 2022).

The goal of this project will be to leverage newly developed RNA foundation models to identify RNAs with potential anti-cancer properties. Predictions of novel anti-cancer therapies will be tested in collaboration with Dr. Rory Johnson (https://www.gold-lab.org/

Keywords

Bioinformatics, therapeutics, machine learning, generative AI

11. QROUTE: An efficient Capacity-Aware Routing Scheme for Quantum Communication Networks

Quantum communication networks promise unprecedented advancements in secure and efficient information transfer by leveraging the principles of quantum mechanics. As these networks evolve, the need for robust and intelligent routing algorithms becomes imperative to ensure optimal resource utilization and information transfer. This doctoral research introduces a ground-breaking Quantum Channel Capacity Aware Routing Algorithm (QCCARA) designed to enhance the performance of quantum communication networks by considering the unique characteristics of quantum channels.

Background: Traditional communication networks face challenges in terms of security and scalability, prompting the exploration of quantum communication as a viable alternative. Quantum communication exploits the properties of quantum entanglement and superposition to secure information transfer against eavesdropping, a feat unattainable in classical systems. However, the efficient routing of quantum information in a quantum communication network poses significant challenges due to the delicate nature of quantum states and the constraints of quantum channels.

Objective: The primary objective of this research is to develop a routing algorithm that optimally utilizes the quantum channel capacity, considering the quantum entanglement and superposition phenomena. By doing so, the algorithm aims to enhance the overall efficiency, security, and reliability of quantum communication networks.

Methodology: The QCCARA is designed based on a comprehensive analysis of quantum channel characteristics, including channel capacity, noise, and entanglement. The algorithm employs advanced quantum error correction techniques to mitigate the impact of channel noise on information transfer. Quantum entanglement is leveraged to establish secure communication channels, and the routing decisions are dynamically adjusted based on the real-time assessment of the quantum channel capacity.

Key Features:

1. Quantum Channel Capacity Assessment: QCCARA dynamically evaluates the capacity of quantum channels, considering factors such as entanglement and noise, to make informed routing decisions.
2. Adaptive Quantum Error Correction: The algorithm incorporates adaptive error correction mechanisms, optimizing information transfer in the presence of quantum channel noise.
3. Real-time Routing Optimization: QCCARA continually adapts its routing decisions based on the evolving conditions of quantum channels, ensuring efficient and secure information transfer.

Expected Contributions: This research is expected to contribute significantly to the field of quantum communication networks by introducing a novel routing algorithm tailored to the unique features of quantum channels. The QCCARA has the potential to improve the overall efficiency and security of quantum communication networks, making them more viable for practical applications.

Conclusion: As quantum communication networks emerge as a transformative technology, the development of sophisticated routing algorithms is crucial to realizing their full potential. The Quantum Channel Capacity Aware Routing Algorithm presented in this research addresses the challenges specific to quantum channels, providing a foundation for the advancement of secure and efficient quantum communication networks

Keywords

Quantum Channel Capacity, Quantum Channel Modelling, Quantum Error Correction based routing, Quantum Communication, Quantum Networks

12. Multi-Agent Pro-Social Rule Bending

With the rise of fields like autonomous vehicles, elder-care robots, and AI assistants in the medical sector, law and governance, one major concern is that goal-oriented agents can cause considerable harm when operating towards the optimal solution for their respective objectives. The need, therefore, for modifying or training these intelligent agents to follow some explicit ethical rules or make decisions with moral reasoning, has now become increasingly important.

Pro-Social Rule Bending (PSRB) is a human response to decision-making scenarios, which exemplifies a form of Care Ethics. That is, an ethical philosophy that considers the contextual information, where the rules are temporarily bent in order to increase the social good. It is an important characteristic of team-work or organizational behaviour, where there is a reflective consideration of the benefits to the 'Other', prioritized over the typical utility function or a behavioural rule. There have been some attempts to encode this into healthcare robots. However, robots and human individuals frequently interact as a part of a team, and robot-agents need to identify when they can dynamically delegate / take-over tasks to/from their team-mate, in order to safely perform PSRB. This form of dynamic teaming between AI and human is an un-explored aspect of multi-agent systems.

In real world scenarios, the contexts will be non-deterministic and the ethical rightness and goodness might change with many parameters (i.e. time, culture, background). As a result, creating a genuinely adaptive agent, with the requisite amount of autonomy is a challenging task. Research in this area would contribute to two areas: machine implementation of ethics, as well as collaborative decision-making in hybrid multi-agent systems.

Keywords

Machine Ethics

13. Neural Network Architecture for Solving Combinatorial Optimisation Problems

Combinatorial optimisation problems arise naturally in many areas of computer science and other disciplines, such as business analytics, operations research, bioinformatics and electronic commerce. Since many of these optimisation problems are NP-hard, applications typically rely on meta-heuristic frameworks, approximation algorithms and carefully designed heuristics for specific instance classes to solve them efficiently. However, the resultant solutions can be very far from optimal and the development of good algorithms often require significant human effort. The goal of this PhD project is to augment the human ability to design good algorithms and data structures by using machine learning techniques to explore the search space efficiently. Specifically, we would like to explore the design of neural network architectures and representations to solve discrete optimisation problems, such as those arising in the context of graphs and geometry. The representations should ideally generalise the known approximation algorithms for classical problems (such as those based on semi-definite programming).

Our research group has done a lot of work in this area and this PhD project will build on this research.

- Lauri, J. et al. (2023) “Learning fine-grained search space pruning and heuristics for combinatorial optimization,” Journal of heuristics, 29(2–3), pp. 313–347.
- Tayebi, D., Ray, S. and Ajwani, D. (2022) “Learning to Prune Instances of k-median and Related Problems,” in 2022 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX). Philadelphia, PA: Society for Industrial and Applied Mathematics, pp. 184–194.
- Fitzpatrick, J., Ajwani, D. and Carroll, P. (2021) “Learning to sparsify travelling salesman problem instances,” in Integration of Constraint Programming, Artificial Intelligence, and Operations Research. Cham: Springer International Publishing, pp. 410–426
- Fitzpatrick, J., Ajwani, D. and Carroll, P. (2023) “Learning to prune electric vehicle routing problems,” in Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 378–392.

I encourage students with a background in mathematics, physics, engineering or business analytics (in addition to the traditional CS background) to get in touch with me.

Students with a background in algorithms or theoretical computer science are particularly encouraged to apply.

Keywords

Combinatorial Optimisation, Machine learning, Graph Algorithms, Algorithm Engineering, Approximation Algorithms

14. Effective Personalised Learning with Generative AI

Decades ago it was shown that personalised learning (one-to-one tutoring) combined with mastery learning resulted in a two standard deviation improvement in student assessment outcomes. This means that the average student learning with a mastery learning approach in conjunction with one-to-one tutoring support can score higher than 98% of students without these supports. In the years since, this has been replicated in experiments - however, in practice personalised learning is exceptionally difficult to achieve at any reasonable scale due to time and resource constraints. The only feasible way of achieving this to-date has been by leveraging funding that almost all institutions cannot justify.

Generative AI is capable of monitoring student progress and providing personalised feedback at scale, for a very acceptable financial cost. A personalised virtual learning assistant will always be available, and can leverage appropriate data to always know where a student is situated with respect to a course of learning. Such a system will also be well-placed to assess not only the products of learning but the process of applying knowledge to various situations and problems. This opens the door to not only developing - but assessing - critical skills, competencies, and dispositions. This project aims to design, develop, and test a personalised learning system powered by Generative AI. The system will be tailored for introductory programming courses as in this domain, Generative AI has been recently shown to be capable of performing in the top percentiles of real university students on real assessments. More advanced courses will be tackled once success at the introductory programming level has been achieved.

This project seeks to assess and mitigate several challenges including hallucinations and bias that are known to occur when using Generative AI. Additionally, technical capability enough is not sufficient to achieve success. The learning design and UX/UI of the application need to be designed to ensure that the application does not inhibit use or impede progress in order to be successful. Further, it has been shown that almost all aspects of Generative AI, from design to training to testing and alignment - all require substantial human involvement. This system will not be designed to replace human instructors, but to augment and compliment human efforts, allowing instructors to leverage one-to-one personalised tutoring combined with mastery learning while providing them with information on each student's progress so that human intervention can be applied where it is needed and not where it is not. This will also free up instructor time allowing them to attend to students when and where it is required while at the same time monitoring group/class progress at a macro level.

The ideal applicant will be a strong programmer with a solid understanding of computing education/pedagogy as well as proficiency in human computer interaction.

Keywords

Generative AI, personalised learning, artificial intelligence, computing education, mastery learning

15. Generative AI for VR: To create the real Ultimate Room (HoloDeck)

In 1964 Ivan Sutherland proposed the concept of the Ultimate Display, which offered the idea of "a looking glass into a mathematical wonderland." This paper proposed "a room within which the computer can control the

existence of matter. " We may not be able to control matter, but with new generative AI and LLM models this is now possible research area, as this concept inspired the Star Trek Holodeck and has become one long term goals of the Virtual Reality research.

The topic is to explore the idea that instead of creating bespoke applications for every VR experience, is it possible to create a generic application that can create any VR experience possible. Even a year ago this would seem like Science fiction but now appear to be feasible research topic.

This PhD proposal is to explore and answer the research questions generated by the creation of such an application as this will become a whole new real research topic within VR and beyond into XR.

Right now we believe the starting point would be create 3D models on fly using tools like Shap-E , take speech using Whisper , process it using LLama 2 to understand what objects should be in the world, and control Avatar's within it. Stable Diffusion could be used to create the skybox. Critical this could offer a new interface for people to interact with LLM models and perhaps become primary computing paradigm.

Keywords

Virtual Reality, LLM, Generative AI

16. AI-Driven Network Intelligence for OPEN Radio Access Network

The PhD Research Project will focus on evolving landscape of Open Radio Access Networks (Open-RAN), which marks a shift from traditional single-vendor setups to multi-vendor environments. This transformation enables the realization of customized and purpose-specific network control functions, strategically placed near the user for real-time control and decision-making. The core of this approach lies in the deployment of xApps/rApps, which interface with Near/Non Real-time RAN Intelligent Controllers (RICs) to make finely-tuned control decisions, a process that is increasingly reliant on models based on Machine Learning (ML) and Artificial Intelligence (AI) due to the complexity, frequency, and multitude of parameters involved in decision-making. Network intelligence will enable Network Automation minimising manual intervention in controlling complex wireless networks.

This project focuses on two key areas:

Development and Enhancement of ML/AI Models for Network Performance Optimization: Researching and creating new ML/AI models or improving existing ones to enhance network performance optimization in O-RAN settings. This includes designing algorithms for dynamic resource allocation, predictive analytics for network traffic, and real-time response to changing network conditions.

Trust and Collaborative Model Training Among Vendors: Exploring the dynamics of trust between various vendors in the O-RAN ecosystem, especially in the context of collaboratively training ML/AI models. This research area aims to address the challenges and opportunities in sharing data and insights among different entities to improve overall network performance and efficiency.

The applicant must have:

- A Strong background in computer/telecommunications networks;
- Working knowledge of at least one programming language such as python, JavaScript, Go and of Linux operating system;
- Good written and oral communication skills;
- In-depth knowledge in at least one of the following areas is essential:
  - Experience with wireless/optical 5G concepts, e.g., Slicing, SDN, NFV, Open-RAN and relevant skills (OpenFlow, ONOS, P4, Virtualisation, and orchestrators)
  - Experience with Blockchain and Smart Contracts, e.g., (Hyperledger Fabric, Ethereum).
  - Practical experience with network emulation/simulation environments (e.g., MININET, NS3, MATLAB)
  - Theoretical tools/methodologies with application in networks including Game Theory, Machine Learning, Optimization.

Keywords

Intelligent networks, Open-RAN, Machine Learning

17. Novel framework and new Dataset for accurate attack detection in EVCSs

This PhD project aims to protect the Electric Vehicle Charging Stations (EVCSs) ecosystem against security attacks. The EVCS, as an IoT device, cannot be decoupled from the Internet in order to offer comprehensive services to the customers. Unfortunately, this enables a set of cyber-attacks against the whole EVCS ecosystem. The effect is not limited to EVCSs alone, as it extends to the power grid and eCAVs, as end users, equally. All these components of the EVCS ecosystem are subject to different kinds of cyber-attacks that

did not exist before or were harder to exploit against traditional networks. Deploying IDSs to monitor malicious activities is crucial for the security of the EVCS ecosystem. Since the effectiveness of IDSs relies on the quality of the training datasets, there is an urgent demand for available up-to-date real-world datasets that are representative for all the attacks specific to EVCSs. Therefore, this PhD project will develop a new anomaly-based IDS framework to secure EVCSs in a more accurate, efficient and flexible manner. This framework will include a novel deep learning (DL) based technique for intrusion detection and a comprehensive new dataset that considers the unique characteristics of the attacks against EVCSs. This dataset will be the first solution to produce a publicly available attack–specific dataset for the evaluation of IDSs for the EVCS ecosystems in reflecting real-life conditions.

Keywords

Intrusion Detection System, Electric Vehicle Charging Stations, Cybersecurity, Attacks, Deep Learning

18. Detecting re-entrancy, timestamp dependence and integer overflow vulnerabilities based on attention-based residual network model

Smart contracts have been used widely in the blockchain industry and the security of smart contracts has attracted the attention of the blockchain industry due to a large number of cyberattacks based on the vulnerabilities of smart contracts. Currently many Ethereum smart contracts are not open-sourced and only publish their bytecode binary files. Thus a model based on source code of smart contracts will fail to detect the vulnerabilities of smart contracts. Besides, a large number of existing methods can't detect the vulnerabilities automatically. Some deep learning models are used to detect the vulnerabilities of smart contracts, however, the training dataset contain only a small number of smart contracts which might result in a poorly trained model. In order to automatically detect the vulnerabilities of smart contracts whose source codes are not available, an attention-based residual network is proposed to detect re-entrancy vulnerabilities, timestamp dependence vulnerabilities and integer overflow vulnerabilities based on a dataset which contains abundant smart contracts with vulnerabilities. First of all, the bytecode of smart contracts is decompiled to opcodes by using Opcode-tool. Then the opcode of smart contracts is preprocessed and attention-based residual network model is utilized to detect multiple vulnerabilities simultaneously. Accuracy and F1 score are utilized as evaluation metrics of this novel vulnerability detection model on Ethereum smart contracts.

Objectives

Propose a deep learning model that can detect re-entrancy vulnerabilities, timestamp dependence vulnerabilities and integer overflow vulnerabilities of Ethereum smart contracts automatically with high accuracy.

Challenges

Currently some researchers have proposed machine learning and deep learning models to detect the vulnerabilities of smart contracts. How to propose a new model that can achieve better performance of vulnerability detection on smart contracts than the existing models? How to evaluate the performance of the novel model with objective metrics?

Keywords

Smart contracts, neural networks, deep learning, Ethereum, vulnerability detection

19. Energy Efficient Digital Twin for Wireless Networks

Digital Twin (DT) technology holds significant potential in various industries, such as Industry 4.0, aviation, and healthcare, enabling immersive digital experiences and proactive utilization of Artificial Intelligence (AI) for enhanced resilience [1]. This project focuses on developing a wireless network-specific Digital Twin to reduce energy consumption. The approach involves creating a proof-of-concept Digital Twin by emulating network behaviour and implementing Software Defined Networking (SDN) using testbeds like Fed4Fire and emulators like Mininet. To minimize energy consumption, task offloading, distributed computing, and workload scheduling will be used in conjunction with algorithms and protocols based on matching theory, reinforcement learning, object association, and power allocation. Additionally, energy-aware transmission protocols will be designed, incorporating data compression, intelligent routing algorithms, and adaptive power control for optimized data transfer efficiency. The project aims to deliver a comprehensive Digital Twin for wireless networks, offering valuable insights, strategies, and solutions to achieve at least a 10% reduction in energy consumption while ensuring a high level of Quality of Service (QoS). This research has the potential to drive innovation across multiple industries and inspire further academic exploration in the field of Digital Twins.

Research Question: How can the development of a Digital Twin for wireless networks, incorporating techniques like task offloading, distributed computing, workload scheduling, energy-aware transmission protocols, and adaptive power control, provide potential solutions to reduce energy consumption by at least 10% while maintaining high Quality of Service (QoS) levels?

Objectives:

1. Explore the trade-offs between energy consumption and Quality of Service (QoS) in wireless networks that employ twinning technology.
2. Develop twinning techniques that prioritize energy efficiency, taking into account operating frequencies, wireless propagation characteristics, resource allocation, object association, and power allocation.
3. Design algorithms and protocols based on matching theory and deep reinforcement learning to optimize resource allocation, object association, and power allocation in wireless networks with twinning.
4. Implement the proposed techniques, considering energy efficiency, QoS, and other relevant performance metrics.
5. Compare the performance of the proposed techniques with existing solutions and validate them using a proof-of-concept digital twin created for a network on a testbed.

The study investigates the balance between energy consumption and QoS in wireless networks with twinning. Energy-efficient techniques will be developed considering factors like frequencies, propagation, resource allocation, object association, and power allocation. Algorithms and protocols based on matching theory and deep reinforcement learning optimize resource and power allocation. Implementation and evaluation will employ comprehensive metrics, including energy savings, QoS, and performance indicators. By incorporating digital twins, we aim to reduce energy consumption by 10% while considering benefits like improved reliability.

[1] H. Ahmadi, A. Nag, Z. Khar, K. Sayrafian and S. Rahardja, "Networked Twins and Twins of Networks: An Overview on the Relationship Between Digital Twins and 6G," in IEEE Communications Standards Magazine, vol. 5, no. 4, pp. 154-160, December 2021.

Keywords

Security and Networks

20. Efficient Data-Free Tabular Model Extraction in Cybersecurity

Recent advancements in machine learning have encouraged numerous businesses to monetise this technology by serving their proprietary models (T) as commercial services. However, developing T is expensive, including large-scale data collecting and labelling, intensive computing resources, and expert knowledge. For example, Open AI's development of GPT-3 is estimated to cost at least $4 million, whereas fine-tuning BERT on the SQuAD dataset or building an application using GPT-3 API only costs around $3 and $200, respectively. Consequently, these well-trained models and their data are the high-value intellectual property of their legitimate owners, which, in contrast, incentivises attackers to steal these models for financial gains. Recently, Data-free model stealing (DFMS) methods have been developed to steal T without knowledge of underlying data and model. The pioneer in this field [1] demonstrated the feasibility of DFMS attacks. This study employed generative adversarial network (GAN) to synthesise Ds optimised for S to learn the predictive functionality of T by approximating the gradients of T using zeroth-order optimisation. However, tabular data's inherent imbalance also poses challenges, making the generator more susceptible to mode collapse and limiting it to explore the T's feature space. To address this issue, [2, 3] proposed to synthesise distinctive samples with higher class diversity, enabling improved exploration of T's feature space. Besides, to extend the use of S to different data domains, the works [4, 5] used contrastive learning (CL) in the test-time adaptation network. This approach results in more distinctive features and makes S achieve higher accuracy on the test domain. However, existing DFMS methods in the literature face several challenges, such as assuming access to surrogate datasets (Ds) similar to Dt or requiring analytic gradient access during inference, which limits their effectiveness when these assumptions are unmet. Success in these attacks relies on the similarity between Ds and Dt, which is difficult to achieve in the real world. Additionally, these methods often involve millions of queries, making them impractical for online attacks with limited query budgets. Finally, research on DFMS in cybersecurity and healthcare domains is scarce, primarily due to unique domain constraints and the characteristics of tabular data. Motivated by the above knowledge gap, this research aims to develop a novel DFMS technique against tabular models in a black-box, hard-labelled and data-free setting and within a tight query budget. This research focuses on the following research questions: 1. How to train an effective GAN to synthesize Ds that closely matches Dt, which in turn accelerates the DFMS process? 2. How to construct a DFMS scheme that propagates strong signals to train the generator and S from T's hard-labelled predictions without requiring its gradients while minimising the number of queries? 3. As the use of S could be extended to a different data domain with a distinct distribution from both the Ds and Dt, how to address the domain shift issue without any of these source data?

References:
[1] 10.1109/CVPR46437.2021.00474
[2] 10.48550/arXiv.2105.08584
[3] 10.1016/j.imavis.2023.104627
[4] 10.48550/arXiv.1901.00976
[5] 10.48550/arXiv.2204.10377

Keywords

AI Security, Adversarial-AI, Model-Extraction, Deep-Learning, Data-Free-Model Stealing, Domain-Adaptation

21. Adaptive Privacy-preserving Techniques for Machine Learning

Recently, during the COVID-19 pandemic, many countries, such as the USA, announced a state of emergency, making it possible to disclose patients' data without authorisation. The HIPAA laws and regulations were suddenly updated after the COVID-19 outbreak, although many individuals' sensitive data might be revealed. Therefore, it is essential to consider the users' awareness of privacy issues and how they can control the disclosure of their sensitive data, ensuring they are leading actors in protecting their privacy. The mere use of a predetermined privacy-preserving machine learning technique such as homomorphic encryption, secure multi-party computation, differential privacy, and trusted execution environments has been proven inefficient in different ways, mainly a trade-off between privacy and the utility of a machine learning model. For instance, when using differential privacy, an incautious noise addition might cause a data distortion, reducing its utility. Some adaptive differential privacy approaches have been proposed to automatically predict individuals' privacy concerns based on their personalities and user-generated content to overcome this issue. Other researchers suggested using hybrid solutions to protect the privacy of both data providers and individuals. Although such hybrid approaches provide both computational and output privacy, they are only considered rough solutions, and the problem of personalised privacy still persists. It is worth mentioning that when the data is distributed between multiple parties willing to build a joint machine learning model on their data, each data owner needs to decide the exposure level and, hence, the optimal technique to be used for preserving the privacy of his data depending on different risk levels.

Challenges

- In heterogeneous environments, it is quite difficult to determine the possible risk for individual privacy leakage.
- It is a big challenge to propose effective privacy-preserving techniques to protect the privacy of individuals. At the same time, diverse types of data and multi-dimensional data are located in different locations.
- How to perform high-performance data analysis in a federated learning environment while continuing to protect user privacy.

Research aim and objectives

The project aims to build adaptive privacy-preserving techniques for machine learning environments that automatically adjust based on privacy risks and data sensitivity levels. At the same time, users are given a chance to decide the optimal privacy solution.

The specific research objectives are

- To investigate and analyse current privacy-preserving techniques in the machine learning environment to build a solid background on the state-of-art of these techniques.
- To investigate the user privacy requirements in the machine learning environment to compute the privacy risk.
- To evaluate the privacy risk level of the users to increase user awareness about their privacy.
- To develop adaptive methods to adjust the optimal privacy-preserving techniques automatically based on the privacy risk level.

Keywords

Privacy, Machine learning, Homomorphic encryption

22. Blockchain-enabled Open-RAN Framework for trusted asset management and supply chain verification

Open Radio Access Network (O-RAN) unlocks the existing closed RAN architecture to open interfaces and protocols to adhere with the market comprising diverse vendors and network softwarization concepts introduced for 5G and 6G. O-RAN, induces new security issues and challenges that is worth investigating to realise the full potential and limitation of the O-RAN. Verification of trusted asset, owner, and supplier identities along with their details, credentials, and security properties in a supply chain within the O-RAN system is vital and can be achieved through Blockchain to ensure distributed trust enhancements.

Objectives

- Design a low-cost and energy efficient blockchain platform for O-RAN tailored to 6G consensus mechanism.
- Enable automated and zero-touch secure O-RAN asset management services using AI-enabled self-learning smart contracts.
- Deploy an automated and zero-touch blockchain based marketplace to trade the automated and O-RAN asset.
- Develop a blockchain based reputation system for all the O-RAN stakeholders to identify the malicious stakeholders.

Expected Results

The developed novel Blockchain enabled O-RAN framework and its applications will be published in reputed journals while its AI-enabled Zero-touch security automation mechanisms will be patented.

Keywords

ORAN, Blockchain, 6G, Security

23. Detection of Adversarial Attacks in AI-based ZSM architectures using Explainable AI

The ZSM (Zero-touch network and Service Management) is the next leap of network management in the telecommunication industry. One of the critical enablers in full E2E automation in ZSM is AI with the help of ML and big data analytics. With the use of explainable AI, there is a great potential to enhance the security of the models. Although there are new explainable AI techniques proposed in the field of computer vision for defence against adversarial attacks, the research in the area of network management is a requirement that needs immediate attention. Automated adversarial sample generation along with XAI techniques for early identification and mitigation is a critical research gap that needs to be addressed urgently.

Objectives

- Developing an automated XAI-based security architecture for detecting adversarial attacks in AI algorithms used for ZSM.
- Investigating current literature and implementations on adversarial attacks and proposed techniques for mitigating them.
- Exploring AI techniques envisaged in ZSM and mapping XAI solutions that are viable candidates for adversarial defence mechanisms.
- Analysis of the performance of virtualised XAI security solutions while exploiting the trade-off between accuracy and interpretability.

Expected results

Development of a robust XAI based security architecture for defence against adversarial attack in ZSM environments.

Keywords

XAI, Security, Adversarial, ZSM, 6G

24. Data synthesis and privacy in 6G

As data underlies most of the applications (especially, privacy-sensitive data) in 6G scenarios, we focus on how to ensure durable, untamperable data ownership to ensure that control of the data usage, as well as liability falls to the right parties. The possibilities revolve around data "watermarking," ideally with a low destruction/tampering footprint on the actual data, so as to retain maximum information. Data synthetization is a prime research candidate to achieve this goal, and current efforts (e.g., MIT's Synthetic Data Vault) are promising but lack addressing more types of data, in a more generic way. For the recipients of third-party data, an important question is whether it is possible to distinguish between synthesized and natural data, as it can impact how much one would trust the AI models trained on the data.

Objectives

- Explore methods for introducing "watermarks" and self-destructive mechanisms into data
- Formulate algorithms for data synthesis in the context of 6G
- Define metrics for evaluating the quality of synthesized data
- Propose and validate detection mechanisms for differentiating external synthesized and natural data

Expected results

We expect the DC to; Publish in journals and high-quality conference publications about the fundamental and experimental results; Produce proof of concepts and demo code and models for DC findings; Produce Patents and intellectual property.

Keywords

6G, Privacy, AI

25. Breaking Encryption from the Outside In

With an increasing mainstream focus on the security and privacy of user data, built-in encryption is becoming commonplace in consumer-level computing devices. Under these circumstances, a significant challenge is presented to the lawful digital forensic investigations where data from encrypted devices needs to be analysed. This project will explore the use of electromagnetic side-channel analysis (EM-SCA) for the purpose of assisting digital forensic investigations on encrypted data devices. EM side-channel analysis is a technique where unintentional electromagnetic emissions are used for externally eavesdropping on the operations and data handling of computing devices. The project aims to develop novel techniques for automated, uninstrumented, AI-driven cryptographic key recovery from IoT devices. While key recovery have achieved in the past using EM-SCA, black box key recovery remains an open issue.

The non-intrusive nature of EM side-channel approaches makes it a viable option to assist digital forensic investigations as these attacks require, and must result in, no modification to the target device. This project involves the combination of a range of technologies including software defined radios, digital forensics and cybersecurity, internet-of-things devices, and deep learning/data science.

Keywords

Digital Forensics, Cybersecurity, Data Science, Artificial Intelligence

26. On the Use of Large Language Models in Security Requirements Traceability

In requirements engineering, traceability is about understanding how high-level requirements – objectives, goals, aims, aspirations, expectations, and business needs – are transformed into development-ready, low-level requirements. It is, therefore, primarily concerned with satisfying relationships between layers of information (aka artifacts). However, traceability may document relationships between many development artifacts, such as requirements, specification statements, designs, tests, models and developed components [1]. Tracing security requirements to code can have the benefits of reducing vulnerability and the effort required to fix them [2]. However, on the one hand software engineers have limited time to implement security functionalities and create traceability links between the software system and the assets to be protected, security goals and requirements. Although several approaches [3, 4] were proposed to perform security requirements traceability, limited research was conducted to suggest automated approaches to trace security goals, requirements and concerns (e.g., assets) in the source code. The advent of Large Language Models (LLMs), such as GPT and Google Bard, has opened up new possibilities for enhancing requirements traceability by leveraging natural language understanding and generation capabilities.

This PhD project aims to explore and develop novel methodologies to utilise LLMs to advance the field of security requirements traceability. We will investigate how LLMs can help software engineers trace assets, security goals and requirements by analysing the source code. We will explore zero-shot and few-shot learning for this purpose and compare the results obtained with LLMs with those obtained by software engineers.

Research Objectives:

- Investigate the state-of-the-art in requirements traceability, LLMs and prompt engineering;
- Elicit tasks that LLMs can perform to annotate security goals, requirements and source code.
- Create and share a dataset with existing open-source software projects annotated with information about security goals, requirements and assets.
- Evaluate the effectiveness of LLM-enhanced techniques in comparison to traditional methods used to perform security requirements traceability.

The candidate PhD student will be part of the Security, Privacy, Adaptation and Requirements Engineering research group (spare.lero.ie) and work collaboratively with other researchers and graduate students in software engineering and security. The prospective student will also be part of Lero - the SFI research centre for Software - (lero.ie) and will avail of the training and outreach opportunities offered by the centre

[1] Cleland-Huang, J., et al., Software traceability: trends and future directions, in Future of software engineering proceedings. 2014. p. 55-69.
[2] Wang, W., et al., Detecting software security vulnerabilities via requirements dependency analysis. IEEE Transactions on Software Engineering, 2020. 48(5): p. 1665-1675.
[3] Breaux, T.D. and D.G. Gordon. Regulatory requirements traceability and analysis using semi-formal specifications. in Requirements Engineering: Foundation for Software Quality: 19th International Working Conference, REFSQ 2013, Essen, Germany, April 8-11, 2013. Proceedings 19. 2013. Springer.
[4] Houmb, S.H., et al., Eliciting security requirements and tracing them to design: an integration of Common Criteria, heuristics, and UMLsec. Requirements Engineering, 2010. 15: p. 63-93.

Keywords

Intelligent networks, Open-RAN, Machine Learning

27. EcoMEC: Optimizing Performance and Sustainability in Edge Cloud Environments

The escalating energy consumption and carbon emissions associated with edge cloud environments present a pressing challenge to sustainable computing. As smart cities increasingly rely on Multi-Access Edge Computing (MEC) services, innovative solutions that balance performance and environmental impact become paramount. The EcoMEC project seeks to address this urgent challenge by comprehensively exploring key research questions. The EcoMEC project aims to pioneer a sustainable computing framework for smart cities to optimize energy usage, mitigate carbon emissions, and integrate renewable energy sources. By achieving this objective, EcoMEC aims to significantly contribute to the field of sustainable computing. To realize this objective, EcoMEC will investigate the following research questions:

1. How can the carbon footprint of edge servers be accurately assessed and compared, laying the foundation for sustainable smart city computing?
2. In what ways can AI be harnessed to promote a green MEC environment for smart city applications, fostering efficiency and environmental responsibility?
3. How can AI-based techniques facilitate elastic resource provisioning in a geo-distributed edge server infrastructure across a smart city, ensuring optimal performance?
4. What scheduling mechanisms can be employed to optimize the allocation of tasks on the geo-distributed edge server infrastructure, considering diverse application requirements and resource constraints within a smart city environment?

The EcoMEC research project acknowledges several challenges and aims to propose innovative solutions to each:

1. Diverse QoS Requirements: Innovations in resource allocation strategies to address diverse Quality of Service (QoS) needs.
2. Dynamic Resource Consumption: Techniques for energy-efficient and carbon-aware resource allocation.
3. Varied Energy Sources: Adaptive mechanisms for managing energy usage amidst diverse renewable and fossil-based sources.
4. Temporal Variation in Energy Source: Intelligent orchestration mechanisms to adapt to the time-variant nature of energy sources.
5. Dynamic Population Changes: Strategies for predicting and adapting to population fluctuations, ensuring optimal resource usage.

The EcoMEC project adopts a mixed-methods approach, combining simulation modeling and empirical analysis. This approach is chosen to comprehensively understand the proposed mechanisms' performance under varying conditions in smart city scenarios. The EcoMEC project's trade-offs and quantitative analysis, contributing to academic knowledge, will also inform real-world decision-making, offering practical insights for smart city planning, resource management, and policy formulation. Aligned with UN Sustainable Development Goal 7, EcoMEC aspires to be a catalyst for advancing sustainable computing practices in smart cities, significantly contributing to global efforts to increase the share of renewable energy in the global energy mix through optimized energy usage, mitigated carbon emissions, and integrated renewable energy sources.

Keywords

MEC, resource management, task scheduling, Carbon footprint

28. GIS-based carbon footprint analysis for urban planning

As cities grow in size, the environmental impact of urbanization is becoming increasingly evident, with escalating carbon emissions and ecological strain. The need for sustainable urban planning that takes into account the carbon footprint of urban activities and structures is important. Traditional approaches to city development often overlook environmental repercussions, resulting in inefficient resource use and heightened impact on climate change. To navigate this challenge, there is a need for a thorough Geographic Information System (GIS)-based analysis of carbon footprints and accompanying tools. Such an analysis platform would serve as a vital tool for urban planners, helping them in making decisions that foster greater sustainable development.

There are a series of technical challenges which need to be considered. At the forefront is the task of (DATA) integrating heterogeneous data at different spatial and temporal scales and ensuring the accuracy and validity of the data from various sectors like transportation, energy, waste, satellite imagery, and land use as well as environmental data. This has not been satisfactorily achieved to date. (MODEL) The creation of a GIS-based model that authentically represents the interplay of direct and indirect emission sources is also a challenge as is demonstrating the generalisability of the model to different regions. (VISUALISATION) As the complexity of data and models increases, presenting the findings in a clear, accessible, and meaningful manner for the target audience becomes crucial for effective decision-making.

Keywords

Geographic Information Science, Spatial Analysis

29.Using ML to enhance the veracity of Open Street Map

Open Street Map (OSM) is a spatial database created and maintained by volunteers. It is well known that there can be errors in the dataset. This PhD topic aims to enhance the accuracy of OSM data using machine learning (ML). OSM, which is a widely-used mapping platform (e.g. Meta), sometimes faces data inaccuracies due to user errors and malicious content. By incorporating ML algorithms, the project aims to automatically detect and rectify these inaccuracies. The research should develop models trained on diverse data sources (internal and external to OSM) to analyze and correct errors, ensuring a more reliable OSM dataset. The goal is not only to fix existing issues but also to create a system that continuously learns and adapts to changes in map data, ultimately improving navigation systems, urban planning, and disaster response capabilities. The research advances the use of machine learning for geospatial data accuracy and can be applied to other sources of spatial data.

The technical challenges for this project revolve around the nature of geospatial data. Searching for outliers and errors is more complex due to the dependent nature of the variables in the data. The OSM dataset is large and complex, determining machine learning algorithms that effectively manage the diversity in the data, accurately identify and categorize errors, and adapt to the dynamic updates in teh ever changing map data is a challenge. In addition to the spatial data, data about the volunteers could be an indicator or data quality and errors. Incorporating this dataset poses additional technical challenges for data fusion. Given the ethos and volunteered nature of OSM, maintaining the Human-in-the-loop is required along with interpretable results. Expertise in machine learning, geospatial data analysis, and an in-depth understanding of the OSM structures and user interactions is required.

Keywords

Machine Learning, Geospatial Data, Data Quality

30. Applying Human Browsing for realising Autonomous Behaviour on the Machine-Readable Web

The idea of a Machine Readable Web orginates from Tim Berner-Lee's keynote at the first Web conference in 1994 [1]. Since this time, a number of technologies have emerged that aim to support this vision, starting with the development of the Resource Definition Framework (RDF) [2], followed by the concept of Linked Data [3] and more recently the emerging concept of Distributed Knowledge Graphs (https://cost-dkg.eu).

From the onset of what has become known as the Semantic Web, it was envisioned that Intelligent Agents and Multi-Agent Systems (MAS) would play a critical role in delivering autonomous behaviour on the web [1]. Unfortunately, due to the lack of tools and testbeds in the early stages of Semantic Web research this never happened [4, 5]. This has now changed, and Semantic Web technologies are increasingly being applied in real world scenarios, most commonly in the form of enterprise knowledge graphs [6]. However, the integration of autonomous behaviour is still an issue that the recently proposed Hypermedia Multi-Agent Systems (MAS) research area aims to address [7]. The main objective of the project is to contribute to this emerging research area by exploring how human browsing tools and behaviour [8,9] can be used to create a browsing model for Hypermedia Agents on the Machine Readable Web. The design of the framework should be sufficiently generic to support the application of Machine Learning techniques. The project will involve developing a conceptual framework of autonomous browsing that will be implemented and evaluated through one or more case studies.

[1] Berners-Lee, Tim. "W3 Future Directions, Keynote." 1st W3 Conference. 1994.
[2] Decker, Stefan, et al. "The semantic web: The roles of XML and RDF." IEEE Internet computing 4.5 (2000): 63-73.
[3] Bizer, Christian, et al. "Linked data on the web (LDOW2008)." Proceedings of the 17th international conference on World Wide Web. 2008.
[4] Hendler, Jim. "Where Are All the Intelligent Agents?," in IEEE Intelligent Systems, vol. 22, no. 3, pp. 2-3, May-June 2007, doi: 10.1109/MIS.2007.62.
[5] McBurney, Peter et al. "The Agents Are All Busy Doing Stuff!," in IEEE Intelligent Systems, vol. 22, no. 4, pp. 6-7, July-Aug. 2007, doi: 10.1109/MIS.2007.77.
[6] Hogan, Adrian et al. “Knowledge Graphs”, url: http://kgbook.org/
[7] Ciortea, Andrei et al. "A decade in hindsight: the missing bridge between multi-agent systems and the world wide web." In AAMAS 2019-18th International Conference on Autonomous Agents and Multiagent Systems, 2019.
[8] Choo, Chun Wei, Brian Detlor, and Don Turnbull. "Information Seeking on the Web--An Integrated Model of Browsing and Searching." (1999).
[9] Agarwal, Naresh Kumar. "Integrating models and integrated models: towards a unified model of information seeking behaviour." Information Research 27.1 (2022): 27-1

Keywords

Hypermedia MAS, Web Architecture, Linked Data

31. Interpretable and Robust Machine Learning Models with Abstention

This project addresses responsible AI by studying robust machine learning models designed to refrain from predictions when uncertain, providing transparent explanations for the abstention. The dual purpose of these explanations is to a) enhance user comprehension and b) empower domain experts to pinpoint inaccuracies for model refinement. As the "uncertainty" part can be a result of noisy data, outliers, overlapping decision boundaries, etc. they are subjective to the domain, therefore ideally this project will validate the outcome of the explanation using user studies. Our focus involves assessing current algorithms, adapting them to this context, or innovating new ones to align with our objectives. Rigorous testing across different domains, ideally incorporating user testing, will gauge the efficacy of the explanations, ensuring the models meet both transparency and functionality standards.

This project is funded through the Beijing-Dublin International College (BDIC). As part of the scholarship, the successful candidate may be expected to travel with the supervisor to Beijing University of Technology in Beijing.

Keywords

Machine Learning, Robustness, XAI, Responsible AI

32. Edge-Embedded AI: Case Study in Healthcare

Edge AI is an emerging trend and is becoming increasingly popular in various real-world applications such as healthcare, industrial automation, smart home devices, and autonomous vehicles. It is characterized by real-time processing and continuous data collection through various devices such as smartphones, IoT, and edge servers. The Edge AI model offers several advantages, including low latency, high data privacy and security, bandwidth, and energy efficiency. Although Edge AI seems very powerful and promising and has numerous advantages, the future of computing is expected to be a hybrid model where all devices coexist to form an extreme-scale computing system. This system will be ideal for edge AI, massive computational power, large-scale data analysis, and collaborative processing across multiple devices (nodes). Federated learning, decentralized and peer-to-peer learning, multi-task learning, and meta-learning approaches are well-suited to run on this computing system. However, this computing environment also presents many challenges. AI in the Edge often involves a combination of algorithmic optimizations, model compression techniques, hardware accelerators, and careful consideration of the specific use case and deployment environment. As edge computing and AI technologies continue to advance, ongoing research and development efforts aim to overcome these challenges.

Keywords

Edge computing, Distributed AI, Mobile Edge, Federated Learning, Mobile data Privacy, Task Offloading