FACULTY OF ENGINEERING AND ARCHITECTURE

IDLab-DECIDE Thesis topics 2025

You can find the full proposals and the point of contact for each topic at our website:

https://decide.ugent.be/

GHENT UNIVERSITY

The DECIDE research group

DECIDE conducts cutting-edge research on Distributed and Embodied Computing In Dynamic Environments.

We focus on AI approaches for:

Robotic Planning

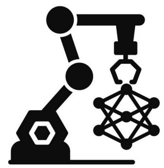

Cognitive control

Distributed perception

Resource efficient AI

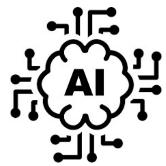

Applied in different domains:

Remote sensing

Precision agriculture

Collaborative robots

Social VR

Industrial robotics

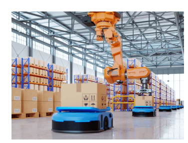

Intelligent environments

Robotic Planning

We develop Generative AI techniques to build AI agents that can imagine the future and plan their actions accordingly.

A (shared) multi-agent policy for dual robot collaboration
Learning a single, overarching policy, or two, "communicating" policies for robotic collaboration.

Autonomous robot navigation in dynamic mazes
Realize a computationally efficient, yet not too greedy navigation strategy in a continuously changing world.

Multi-agent RL in a small-scale robot town
Develop RL policies for dynamic task allocation and execution between mobile robots.

IDLab INTERNET & DATA LAB

Cognitive Control

We model human decision-making patterns and social interactions to build seamless human-AI collaboration.

Teacher-student scenario with simulated students using Large Language Models
Develop a framework to train a teacher using simulated students in VR.

Progress tracking in virtual reality
Automatic advice generation using LLMs for users based on their progress in VR applications.

Imitation learning from low-quality demonstrations
Develop imitation learning algorithms that mimic expert behavior and can deal with imperfect demonstrations, gaps, or dynamics mismatches .

Cooperation between RL agents with non-exponential reward discounting
Investigate how cooperative multi-agent RL (MARL) dynamics change when moving from agents with exponential discounting to human-like non-exponential discounting.

Distributed Perception

We explore how AI techniques can be used in combination with sensors such as cameras, LiDAR and Hyperspectral Imaging in distributed settings.

Hyperspectral super resolution
Use AI techniques to generate images with a high spatial and spectral resolution by fusing RGB and hyperspectral data.

Mine detection using Ground Penetrating Radar
Develop AI algorithms to detect buried landmines with data from a Ground Penetrating Radar.

Looking beneath the Surface of Roman Republican Cities using Ground Penetrating Radar
Perform anomaly detection on Ground Penetrating Radar data to support archeological surveys of Roman Republican Cities.

Mitigating Privacy Risks in Non-RGB Imaging Sensors Using Deep Neural Networks
Explore the privacy risks associated with non-RGB camera modalities and develop effective solutions that balance privacy concerns with data utility.

Resource efficient AI

We investigate approaches to reduce the computational cost of deep neural networks to make them compatible with resource constrained devices.

AI algorithms for novel hardware
Map combinatorial and planning problems onto an Ising machine consisting of a lattice of probabilistic bits.

Echo state networks for object detection
Explore the potential of ESNs for object recognition, focusing on their ability to perform efficient feature extraction while maintaining competitive accuracy.

Knowledge Distillation for Vision Transformers
Develop novel knowledge distillation approaches to improve the efficiency of large Vision Transformer models.

Chaos Computing for Deep Neural Networks Optimizations
Develop a chaos-driven adaptive pruning framework for deep neural networks that reduces the computational cost at runtime based on the complexity of the input data.

Efficient AI processing onboard satellites
Optimize deep neural networks for onboard processing on satellites.

GHENT UNIVERSITY

IDLab
INTERNET & DATA LAB

Efficient AI processing onboard satellites

Contact: Sam Leroux: sam.leroux@ugent.be

Problem statement

Satellites are increasingly being integrated with onboard AI algorithms, enabling them to perform complex tasks such as real-time data analysis, autonomous navigation, and anomaly detection. This combination enhances the efficiency and accuracy of satellite operations, opening up new use cases in areas like environmental monitoring, disaster response, defense and precision agriculture. However, deploying AI in space presents challenges, including limited computing resources, data transmission delays and the harsh space environment.

Thesis goals

This thesis is in collaboration with the Ghent-based company EDGX. The goal of this thesis is to develop novel efficient AI applications suitable for deployment on the custom AI computing payload developed by EDGX. Depending on the interest of the student, the focus can be on different applications such as satellite communication, earth observation or in-orbit servicing.

Chaos Computing for Deep Neural Network Optimizations

Contact: Vishisht Sharma: vishisht.sharma@ugent.be , Ciem Cornelissen: ciem.cornelissen@ugent.be

Problem statement

Traditional deep learning models often suffer from excessive computational overhead due to static network architectures that do not adapt to varying input complexities. This inefficiency is particularly evident in image processing tasks, where different images require different levels of computational effort. Current neural pruning techniques lack dynamic adaptability, leading to either excessive resource consumption or loss of critical information. This research aims to explore applications of chaos-based techniques for optimization. For example, a **chaos-driven adaptive neural pruning framework** leverages chaotic computing principles—such as logistic maps—to dynamically adjust the network's computational complexity based on the complexity of the input images. By intelligently pruning and activating neural connections in real time, this approach seeks to enhance computational efficiency while maintaining high accuracy across diverse image datasets.

https://arxiv.org/abs/2308.09955

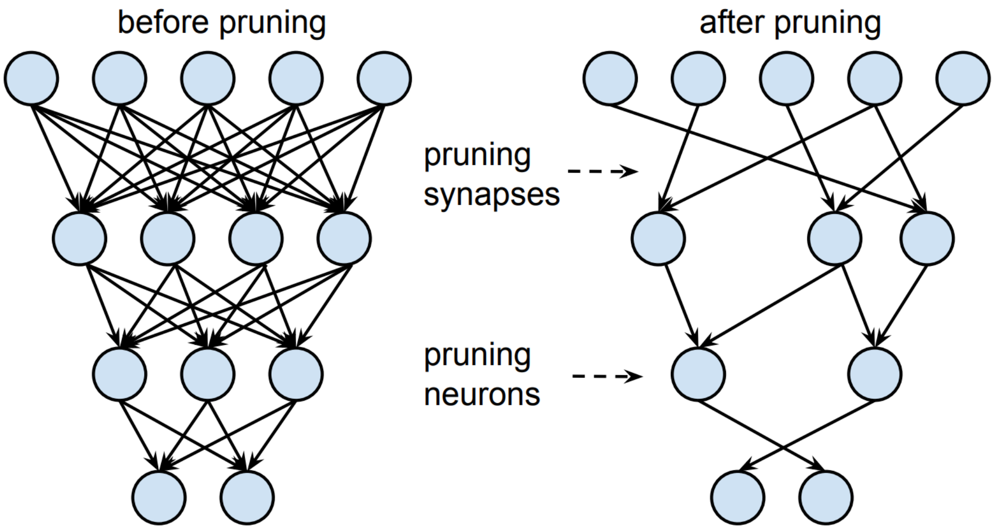

Thesis goals

One example of such research is to develop a chaos-driven adaptive neural pruning technique that dynamically adjusts the computational complexity of deep learning models based on the complexity of images. This approach leverages chaotic computing principles, such as logistic maps, to determine how many and what should be pruned or activated at any given time. Other directions can be explored as well.

Deep Learning Based Compression of Hyperspectral Data

Contact: Ciem Cornelissen: ciem.cornelissen@ugent.be

Problem Statement

Hyperspectral imaging captures high-dimensional data across numerous spectral bands, providing detailed insights into material properties. This makes it invaluable in applications such as precision agriculture, remote sensing, and medical imaging. However, the large size of hyperspectral data poses significant challenges for storage, transmission, and real-time processing, especially in embedded and edge computing environments.

Traditional compression techniques, such as PCA-based or wavelet-based methods, often struggle to maintain spectral fidelity while achieving high compression rates. Recent advances in deep learning have shown promise in learning compact representations that preserve essential information while significantly reducing data size.

This thesis aims to explore deep learning-based methods for compressing hyperspectral data while maintaining critical spectral and spatial information, making it more suitable for efficient storage and real-time applications.

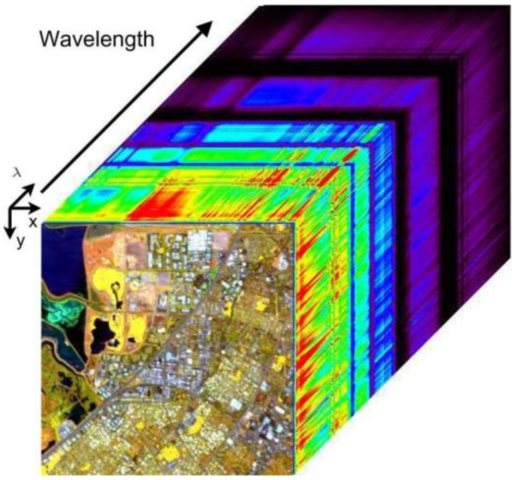

Thesis Goals

The primary goal of this thesis is to develop a deep learning model capable of efficiently compressing hyperspectral images while preserving their key features. Some of the key objectives could be but are not limited to:

• **Developing an autoencoder or transformer-based architecture** to learn compact latent representations of hyperspectral data.
• **Exploring hybrid compression approaches** that combine deep learning with traditional techniques for improved efficiency.
• **Evaluating performance trade-offs** between compression ratio, reconstruction quality, and computational complexity on real-world datasets.
• Optimizing the method for real-time and embedded applications, ensuring feasibility for edge devices

Possible Approach

This thesis will investigate multiple deep learning techniques for hyperspectral data compression, including:

- **Autoencoders & Variational Autoencoders (VAEs):** Learning compact representations while ensuring minimal reconstruction error.
- **Transformers & Attention Mechanisms:** Leveraging spectral and spatial dependencies for efficient compression.
- **Quantization & Pruning:** Reducing model size and computational overhead while maintaining data integrity.
- **Lossy vs. Lossless Compression:** Exploring how different approaches affect hyperspectral image utility for downstream tasks such as classification and segmentation.

Early Confidence Prediction for Adaptive Efficient Deep Learning Inference

Contact: Ciem Cornelissen: ciem.cornelissen@ugent.be , Vishisht Sharma: vishisht.sharma@ugent.be

Problem statement

Deep learning models are widely used in various applications, such as computer vision, and large language models. However, these models often require high computational resources, making them impractical for real-time or embedded deployment. One way to improve efficiency is through **adaptive inference**, where computations are dynamically adjusted based on model confidence.

This thesis explores **early confidence prediction**, a technique that estimates a model's confidence early in the inference process. If confidence is sufficiently high, computations can be **skipped** or an **early exit** can be triggered, reducing processing time while maintaining accuracy. This approach is particularly beneficial for resource-constrained environments such as edge devices and real-time applications.

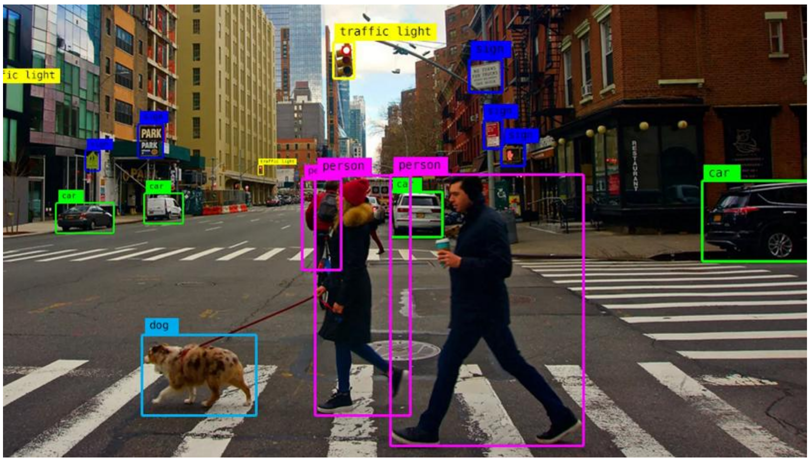

Thesis goals

The goal of this thesis is to develop and evaluate a deep learning method that predicts model confidence early during inference, enabling adaptive efficiency mechanisms. Some of the key objectives could be but are not limited to:

• **Develop an early confidence predictor** that estimates the final model confidence using intermediate feature representations.

• **Implement adaptive inference mechanisms**, such as layer skipping, early exiting, or dynamic resolution adjustment, based on the predicted confidence.

• **Evaluate performance trade-offs** between efficiency and accuracy on real-world datasets, such as hyperspectral imaging, automotive perception, or precision agriculture.

Possible Approach:

This thesis will explore multiple techniques to estimate early confidence and implement adaptive inference:

• **Auxiliary Networks**: Train small sub-networks on intermediate activations to predict confidence scores.

• **Uncertainty Estimation**: Apply techniques such as Monte Carlo Dropout, Bayesian neural networks, or deep ensembles to quantify confidence.
• **Dynamic Skipping Mechanisms**: Implement conditional layer skipping, early exits, or adaptive resolution scaling based on confidence predictions.

Object Recognition with Echo State Networks

Contact: Vishisht Sharma : vishisht.sharma@ugent.be

Problem statement

Object recognition is a fundamental task in computer vision, traditionally dominated by deep learning models such as Convolutional Neural Networks (CNNs). However, these models often require extensive training, high computational power, and large amounts of labeled data, making them inefficient for resource-constrained environments. Echo State Networks (ESNs), a form of Reservoir Computing, offer a promising alternative by leveraging sparse and randomly initialized recurrent networks to efficiently process visual data.

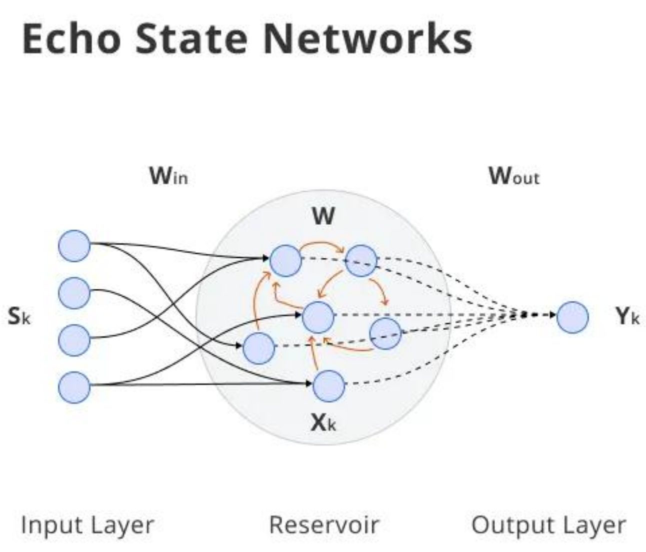

This research aims to explore the potential of ESNs for object recognition, focusing on their ability to perform efficient feature extraction while maintaining competitive accuracy. The study will investigate a multi-reservoir approach, optimize hyperparameters for improved classification, and compare ESN-based models against traditional deep learning techniques in terms of performance, robustness, and computational efficiency.

Thesis goals

The goal of this research is to explore, develop and evaluate an **Echo State Network (ESN)-based approach for object recognition**, leveraging the advantages of reservoir computing for efficient and accurate classification.

Looking beneath the Surface of Roman Republican Cities using Ground Penetrating Radar

Contact: Sam Leroux: sam.leroux@ugent.be

Problem statement

Ground-penetrating radar (GPR) has revolutionized archaeological investigations by allowing researchers to uncover hidden structures and artifacts without excavation. By transmitting electromagnetic waves into the ground and analyzing the reflected signals, GPR can create detailed images of subsurface features, revealing buried walls, foundations, and even entire settlements. This non-invasive method preserves the integrity of archaeological sites while providing invaluable data. However, the sheer volume of data generated by GPR can be overwhelming and time-consuming to process manually. This is where automated AI-based processing steps in, offering advanced algorithms that can quickly and accurately interpret GPR data, highlighting areas of interest and accelerating the discovery process. The integration of AI not only enhances efficiency but also opens new possibilities for uncovering our past with unprecedented precision.

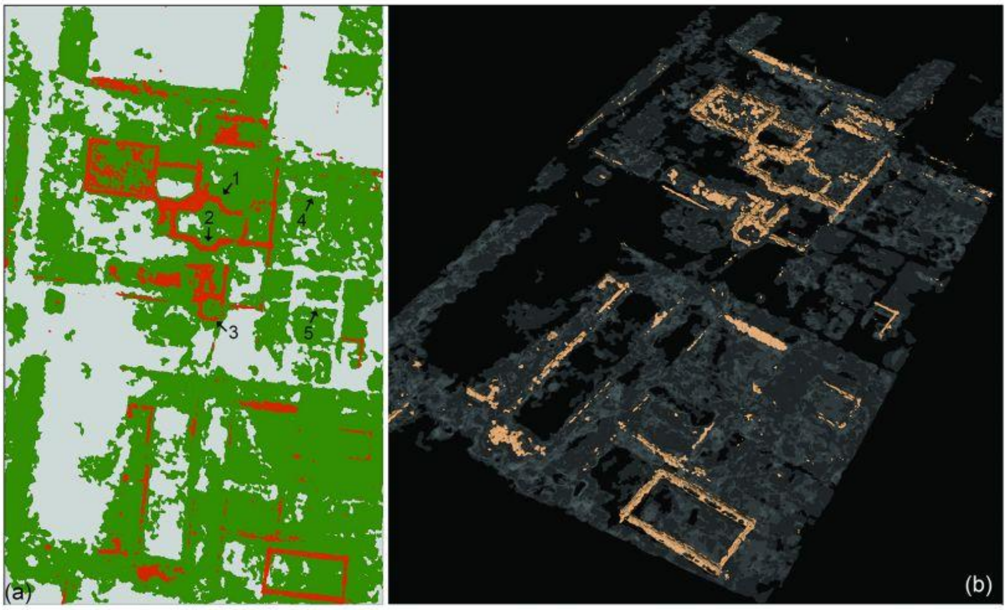

Thesis goals

The goal of this thesis is to develop automated machine learning algorithms to efficiently process GPR data. Depending on the interest of the student the focus can be on anomaly detection, noise reduction or 3D rendering.

Automatic land mine detection using ground penetrating radar

Contact: Sam Leroux: sam.leroux@ugent.be

Problem statement

Despite diplomatic efforts such as the Ottawa Treaty, landmines remain widely in use in conflicts all around the world. In the current Ukrainian-Russian conflict for example, anti-personnel and anti-vehicle mines are used by both parties. An estimated 40% of Ukrainian territory is affected by mines [1]. Even after the conflict would have ended, neutralizing the estimated 2 million landmines [1] would cost over 38 billion USD [2].

Novel unmanned platforms combined with different types of sensors can be used as a safe and scalable alternative to manual demining. Ground Penetrating Radar can be used to "see" below the surface. However, interpreting the output requires a lot of expertise.

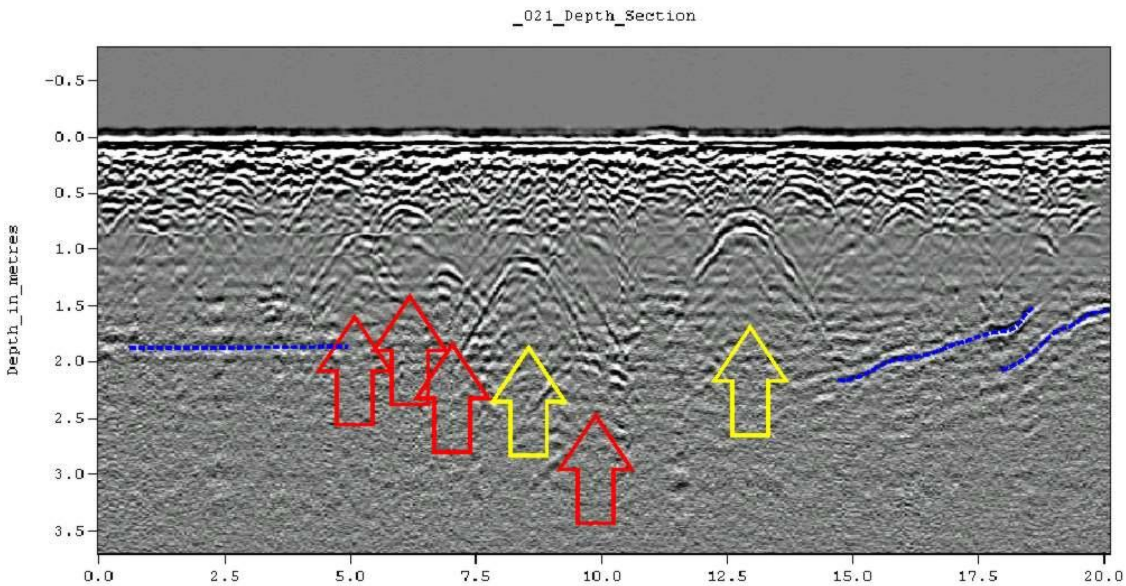

Thesis goals

The goal of this thesis is to develop automated machine learning algorithms to detect mines and other potential underground threats in GPR data.

[1] https://cepa.org/article/an-explosive-choice-landmines-and-ukraine/
[2] https://www.osw.waw.pl/en/publikacje/osw-commentary/2023-11-22/ukraine-worlds-biggest-minefield

Hyperspectral Super Resolution

Contact: Sam Leroux: sam.leroux@ugent.be, Ciem Cornelissen: ciem.cornelissen@ugent.be

Problem statement

Hyperspectral imaging is a powerful technology that captures and processes information across a wide spectrum of wavelengths, beyond what the human eye can perceive. Unlike traditional imaging techniques that use only a few spectral bands, hyperspectral imaging divides the spectrum into numerous narrow bands, allowing for detailed analysis of materials and objects. This technology is widely used in agriculture for crop health monitoring, in environmental science for detecting pollution, and in medicine for diagnosing diseases through non-invasive tissue analysis.

Due to physical limitations, all cameras need to find a trade-off between affordability, sampling rate and spatial and spectral resolution. Hyperspectral Super Resolution aims at automatically generating higher resolution images based on a low-resolution version and would allow to circumvent these limitations.

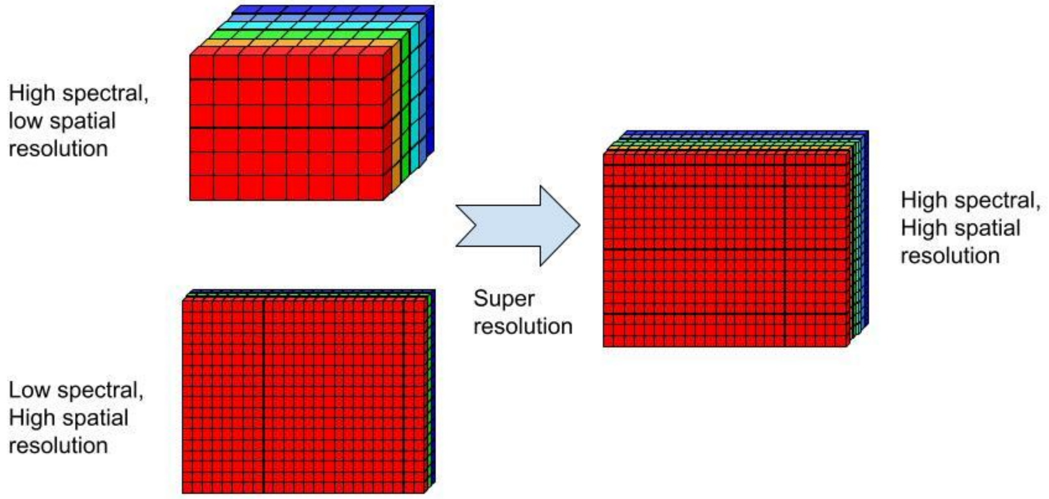

Thesis goals

The goal of this thesis is to develop a new, deep learning based technique for hyperspectral super resolution. Depending on the interest of the student, we might focus more on combining multiple hyperspectral images into one high resolution version or on combining a low resolution hyperspectral image with a high resolution RGB image. This technique will be evaluated on real world data from a precision agriculture or earth observation task. The key focus will be to make this technique suitable for embedded platforms.

Knowledge Distillation for Vision Transformers

Contact: Vishisht Sharma : vishisht.sharma@ugent.be , Ciem Cornelissen: ciem.cornelissen@ugent.be

Problem statement

Vision Transformer (ViT)-based segmentation models have demonstrated exceptional performance in various computer vision tasks. However, their deployment in real-time applications or resource-constrained environments is hindered by their high computational cost and memory requirements. Traditional model compression techniques, such as pruning and quantization, often result in significant performance degradation. Knowledge Distillation offers a promising solution by transferring knowledge from a large teacher model to a smaller student model while preserving accuracy. However, existing distillation approaches for transformers have been mainly focused on LLMs. This research aims to extend KD to vision transformers by exploring and developing effective knowledge distillation techniques tailored for ViT-based models, ensuring efficient performance.

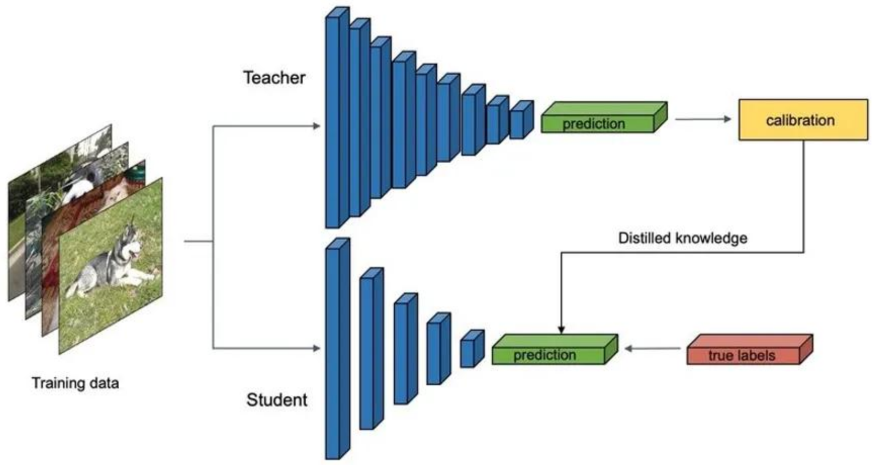

Thesis goals

The goal of this thesis is to develop and optimize knowledge distillation techniques specifically tailored for Vision Transformer (ViT)-based models. By exploring various teacher-student architectures, loss functions, and attention-based distillation strategies, this research aims to enhance model efficiency while preserving high accuracy. The proposed methods will be evaluated on benchmark datasets such as A2D2 and Cityscapes to validate their effectiveness in real-time and resource-constrained environments.

Imitation learning from low-quality demonstrations

Contact: Elias Malomgré: elias.malomgre@ugent.be

Problem statement

Imitation learning is a machine learning approach where an agent learns to perform tasks by mimicking expert demonstrations instead of relying on explicit reward signals. It aims to generalize the demonstrated behavior to new situations by mapping observed states to appropriate actions, often using techniques like behavioral cloning, inverse reinforcement learning, or Generative Adversarial Imitation Learning. However, imperfect demonstrations and demonstrations with gaps or missing actions can severely impact the learning process and the performance of the resulting policy. Imperfect demonstrations may include suboptimal, inconsistent, or noisy actions that mislead the agent, causing it to learn ineffective or unstable behaviors. Gaps or missing actions can disrupt the sequential nature of decision-making, making it difficult for the agent to infer the correct action in certain states. This can lead to incomplete or mistaken policy learning, especially if the missing actions correspond to critical decision points, making it harder for the agent to generalize effectively. Handling such imperfections often requires additional techniques like data augmentation, reward shaping, or integrating human feedback to ensure more robust learning. There can also be a mismatch in the dynamics of the expert environment, which can be different from the training environment. This makes it not trivial to mimic and transfer the behavior of the expert.

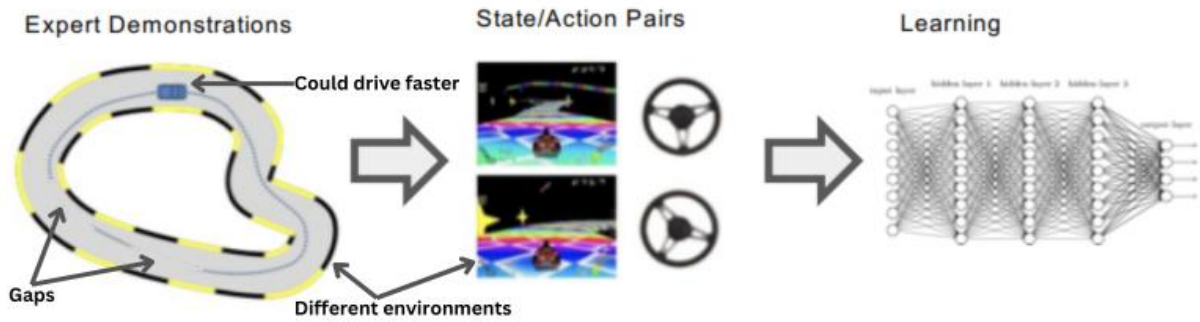

Thesis goals

The aim of the thesis is to address these limitations and develop algorithms that still mimic expert behavior. This topic has been divided into three theses, each dealing with one of the limitations, and the student can choose one of them to research.

- Thesis 1 will focus on developing or extending methods to address gaps in demonstrations and how different environments change the best approach.

- Thesis 2 will tackle the problem of imperfect demonstrations by coming up with techniques to detect demonstrations where the actions are not optimal or augment rewards to overcome the problem.
- Thesis 3 will work on finding ways to mimic and transfer behavior between environments with different dynamics.

Cooperation between RL agents with non-exponential reward discounting

Contact: Pieter Simoens: pieter.simoens@ugent.be

Problem statement

In

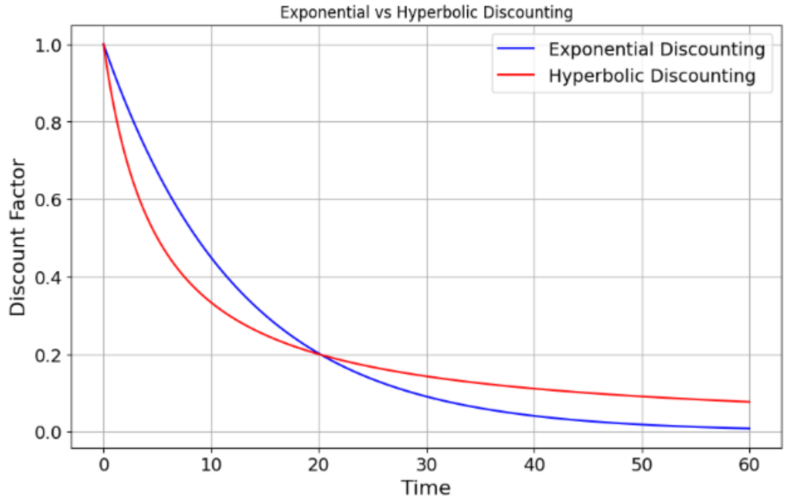

Reinforcement Learning (RL), the objective is to learn a policy that maximizes the cumulative reward resulting from the actions taken. Typically, future rewards are exponentially discounted: a reward that will be obtained n steps in the future is discounted by a factor of γ^n (where γ is the constant discount factor).

However, behavioral research has shown that humans do not discount future rewards exponentially. Instead, they often follow a hyperbolic discounting model, where future rewards are discounted by a factor of 1/(1+n). Hyperbolic discounting leads to time-inconsistent behavior where preferences change over time in a way that leads to suboptimal long-term decision-making.

For example, an individual planning a healthy diet in spring might expect to resist the temptation of summer barbecues. However, when the moment arrives, the immediate temptation of the barbecue may outweigh their original commitment. Similarly, hyperbolic discounting contributes to procrastination, where students defer studying until right before a deadline despite earlier intentions to work consistently.

As we envision a future where AI agents will integrate and cooperate in both machine-machine and human-machine interactions, it is essential to study how such effects affect cooperative behavior in multi-agent settings.

Thesis goals

This thesis aims to investigate how cooperative multi-agent RL (MARL) dynamics change when moving from agents with exponential discounting to non-exponential discounting. We will focus on mixed-motive cooperation scenarios, where agents must navigate both social dilemmas (self-interest vs group benefit) and intertemporal trade-offs (short-term vs long-term rewards). The MeltingPot suite of Deepmind provides a number of mixed-motive environments, but creative students are welcome to propose their own environments.

The first phase of the thesis consists of a comparative analysis of agent behavior under exponential vs non-exponential discounting in cooperative RL tasks. You will adapt existing single-agent hyperbolic discounting techniques to multi-agent RL settings. Based on the results, you will formulate hypotheses about how non-exponential discounting leads to time-inconsistent cooperation strategies.

In the second phase, you will develop self-control mechanisms for hyperbolic RL agents to mitigate time-inconsistent behavior. You can for instance explore pre-commitment strategies, where agents take early costly actions to prevent future undesired behaviors. Humans exhibit pre-commitment behaviors, such as purchasing a long-term costly gym membership to encourage sustained future exercise. You will investigate how self-control mechanisms influence long-term cooperation and fairness in multi-agent environments.

By completing this thesis, you will gain expertise in multi-agent RL, computational modelling of social dilemmas and mechanisms to overcome decision biases in AI systems.

Unseen Threats: Mitigating Privacy Risks in Non-RGB Imaging Sensors Using Deep Neural Networks

Contact: Sander De Coninck (sander.deconinck@ugent.be), Ciem Cornelissen (ciem.cornelissen@ugent.be)

Problem statement

When thinking about the risks of having sensitive information being compromised through cameras, RGB cameras are often pointed to as the main culprit. Emerging alternative imaging sensors, such as Radar, LiDAR and Infrared sensors, are often considered more privacy-preserving alternatives as they capture minimal or no details on colour or texture, rendering tasks like identification through face recognition unfeasible. However, recent research unveils the surprising versatility of these alternative sensors. They can be utilized for a variety of intrusive detections, including emotion recognition, identifying personal attributes (such as height, age, and gender), tracking individuals and even identifying people based on their unique gait [1, 2, 3, 4]. Clearly, relying solely on the replacement of RGB cameras by any of these sensors does not ensure privacy protection. Therefore, there is a need to develop robust methodologies that selectively remove sensitive information while still enabling essential tasks. Imagine a scenario where faces in a video stream are blurred before conducting object classification. A conceptual view of this proposal can be seen in the figure below.

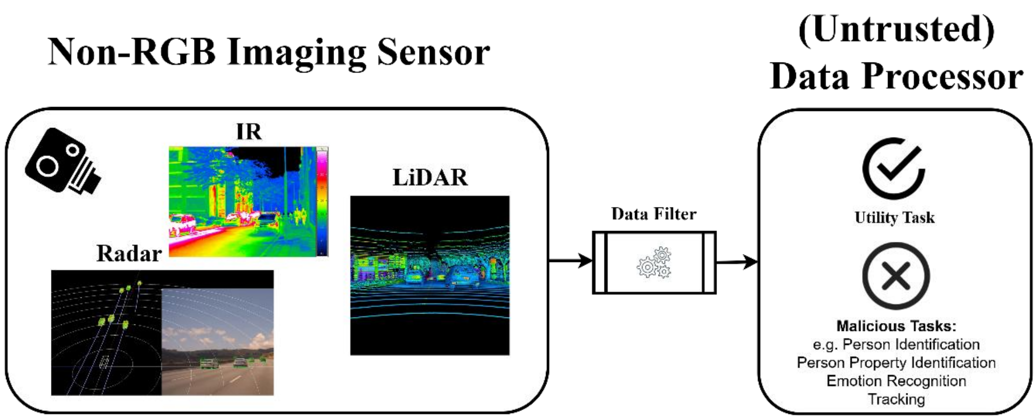

Thesis goals

This thesis aims to thoroughly explore the privacy risks associated with non-RGB camera modalities and develop effective solutions that balance privacy concerns with data utility. Through a literature review, you will explore the diverse range of tasks viable for Radar, Lidar or thermal cameras and have the opportunity to select a modality of interest, setting the stage for the research. In the next phase, you will identify or collect a suitable dataset and define acceptable and prohibited tasks to be performed on this data. Finally, deep learning techniques must be applied to create a tailored filter that, for instance, allows object detection on LiDAR data but prevents person identification.

Throughout this thesis, you will gain practical experience using deep learning techniques on non-RGB camera modalities such as Radar, Lidar or infrared, alongside insights into privacy-preserving machine learning. One possible method involves using generative adversarial privacy, in which generative modelling techniques like GANs or VAEs are used to transform data into a privacy-sensitive version.

[1] Y. Yang, C. Hou, Y. Lang, G. Yue, Y. He, and W. Xiang, “Person identification using micro-doppler signatures of human motions and uwb radar,” IEEE Microwave and Wireless Components Letters, vol. 29, no. 5, pp. 366–368, 2019.
[2] J. Gong, X. Zhang, J. Ren, and Y. Zhang, “The invisible shadow: How security cameras leak private activities,” in Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, 2021, pp. 2780–2793.
[3] M. Hasan, J. Hanawa, R. Goto, R. Suzuki, H. Fukuda, Y. Kuno, and Y. Kobayashi, “Lidar-based detection, tracking, and property estimation: A contemporary review,” Neurocomputing, 2022.
[4] C. Gouveia, A. Tom´e, F. Barros, S. C. Soares, J. Vieira, and P. Pinho, “Study on the usage feasibility of continuous-wave radar for emotion recognition,” Biomedical Signal Processing and Control, vol. 58, p. 101835, 2020.

Progress tracking in virtual reality

Contact: Elias Malomgré: elias.malomgre@ugent.be

Problem statement

Progress tracking is increasingly becoming a critical component in various projects and applications. It allows individuals and teams to monitor their advancement towards set goals, identify potential obstacles, and make necessary adjustments to stay on course. Effective progress tracking not only provides a clear view of what has been accomplished but also highlights areas that require further attention, thereby enhancing productivity and ensuring the timely completion of tasks. The use of VR is increasing rapidly as advancements in technology make VR more accessible and versatile. VR is being adopted across various industries, including gaming, education, healthcare, real estate, and training simulations, allowing users to experience immersive and interactive environments. Tracking progress in VR is essential as it enables users to monitor their interactions within the virtual space, assess their performance, and receive real-time feedback. Virtual environments offer a wealth of metrics that can be leveraged for progress-tracking algorithms, such as elements within the environment and human gestures. Utilizing these metrics, the proposed system aims to comprehensively and accurately represent progress within virtual spaces that can be used as input to LLMs to provide feedback to the user.

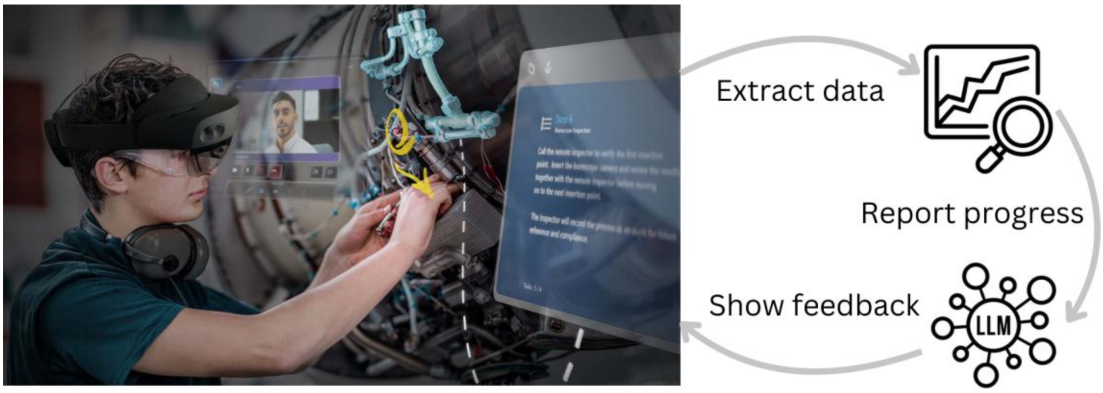

Thesis goals

This thesis aims to develop and implement a robust progress-tracking system for virtual environments. By ents. By utilizing the rich metrics available in VR, such as the interaction with virtual objects and user movements, the proposed system aims to provide precise and real-time updates on user progress that will be used as input to LLMs to generate real-time feedback as guidance for the user. This topic will consist of two different master thesis topics from which the student can choose:

- Topic 1 aims to develop a progress tracking system that can detect at which step in the task the user is with what the user has tried during the step and how long he has been stuck. This information will then be passed to an LLM to generate hints similar to predefined hints for each step.
- Topic 2 focuses on creating a module that will predict which actions are the best to take at each moment by training an agent on the task at hand and using the predicted actions as guidance. This action or sequence of actions will then be given to a LLM to generate hints.

Robotic navigation: autonomous navigation in dynamic mazes

Contact: Daria de Tinguy (daria.detinguy@ugent.be), Matteo Cardoni (matteo.cardoni@ugent.be)

Problem statement

Autonomous navigation in unknown environments remains a big challenge, especially in environments that dynamically change. Existing methods, such as Dynamic Windows Approaches indeed struggle to navigate in these dynamic environments, but also have difficulties in long-term planning. The latter is needed to realize efficient, not too greedy navigation strategies.
AI-inspired strategies on the other hand are very promising in this respect, however at the price of computationally heavy training as well as requiring considerable compute power when effectively navigating.

The challenge of this master dissertation is to explore new (potentially AI-inspired) navigation strategies, able to cope with dynamic mazes.

Thesis goal

The research will start by familiarizing yourself with state-of-the-art obstacle avoidance strategies, including classical approaches (potential fields, Dynamic Window Approach) as well as modern AI-oriented counterparts (mostly based on Deep Reinforcement Learning or variants thereoff). Assessing performance of these algorithms will be through simulations and/or real world deployments in our lab (depending on your interest). After this first phase, opportunities to improve these baselines will be investigated. The end goal is to show case a navigation strategy in a dynamic maze.

Teacher-student scenario with simulated students using Large Language Models

Contact: Elias Malomgré: elias.malomgre@ugent.be

Problem statement

Training an effective teacher to guide a student through a task requires interaction between the teacher and the student. However, relying on a human student throughout the training process is impractical due to time constraints and variability in human behavior. Therefore, a simulated student is needed to enable efficient teacher training. The challenge lies in ensuring that the simulated student accurately reflects the behavior of a real student. Otherwise, the trained teacher may develop strategies that fail in real-world applications due to a mismatch between expected and actual student behavior. This challenge becomes even more complex in virtual reality (VR) environments, where the student's avatar must exhibit realistic movements and interactions. Suppose the simulated student's behavior or movements deviate significantly from a real student's. In that case, the teacher may learn ineffective teaching strategies, ultimately reducing the effectiveness of the trained model when deployed in real scenarios.

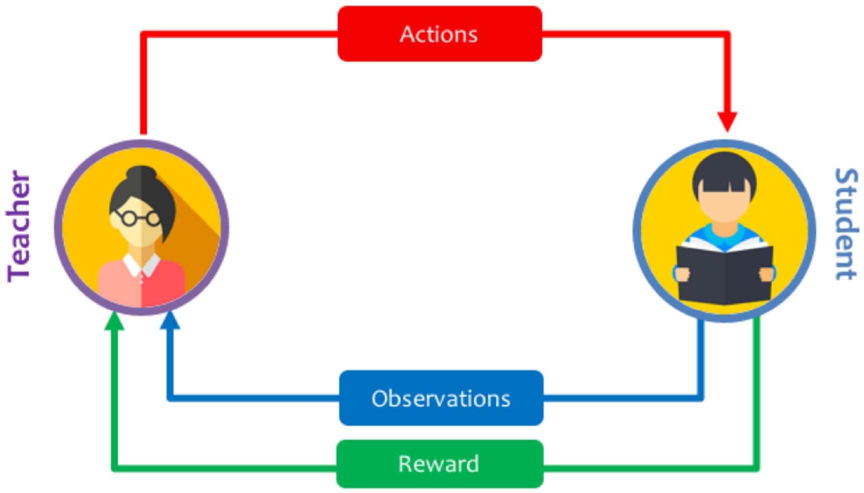

Thesis goals

This research is divided into two complementary theses, each focusing on a critical aspect of the problem to address the challenges of training a teacher through interaction with a simulated student.

The first thesis aims to develop and evaluate a realistic student simulation for teacher training, primarily focusing on modeling student-teacher interactions. This includes designing a student agent that responds naturally to the teacher's guidance, exhibits appropriate learning behavior, and provides meaningful feedback to improve the teacher's strategies. The goal is to ensure the teacher learns effective instructional techniques that generalize well to actual students.

The second thesis focuses on enhancing the realism of student movements in a VR environment. This involves generating naturalistic avatar motions using techniques such as reinforcement learning or large language models to produce lifelike behaviors that align with real human movements. The objective is to create a student simulation that interacts realistically and moves in a way that enhances immersion and usability in VR-based training scenarios.