Project proposal number: 1

Name of the proposed academic supervisor: Christian Berger

Name of the proposed industrial supervisor: Qinglei Ji

Project proposal description:
Modern advanced driver assistance systems (ADAS) like advanced emergency braking systems (AEBS) and automated lane keeping systems (ALKS) aim to reduce the number of fatally injured traffic participants. Systems that enable robotaxis report that AI/ML-enabled systems play a key role in an automated driving system (ADS) to perceive their surroundings. Yet, such ADS inevitably face complex real-world challenges that already resulted in tragic traffic accidents. Out-of-Distribution (OOD) detection approaches can help to evaluate the confidence of such AI/ML-enabled systems and to guide the overall ML development process. However, systematic benchmarking concepts for OOD covering real-world settings with multi-modal, time-series data are scarce to effectively curate OOD datasets that are free from in-distribution samples. In this project, we invite WASP PhD students to systematically explore and analyze the performance of LLMs/FMs to act as OOD detectors on the example of automotive datasets. We plan to use the Zenseact Open Dataset (ZOD) and state-of-the-art multi-modal LLMs like Gemini3 and LLaVA to setup an experimental OOD pipeline that we systematically feed with known in-distribution samples and out-of-distribution samples to investigate modern GenAI's performance to serve as OOD guard.

Competence keywords: * MLops

* LLMs

* Foundation Models

* Python

* Curating datasets

Connection to WARAs: None

Project proposal number: 2

Name of the proposed academic supervisor: Xuan-Son (Sonny) Vu

Name of the proposed industrial supervisor: Assigned by NVIDIA Higher Education, tentatively Zenodia Charpy

Project proposal description:
This project focuses on the development of advanced Hybrid Retrieval-Augmented Generation (RAG) systems to enhance the performance of intelligent agents in long-term interactions. While traditional RAG systems primarily rely on semantic embeddings for retrieving relevant information, Hybrid RAG integrates structured knowledge graphs alongside vector-based retrieval. This approach improves accuracy, contextual relevance, and interpretability.

LLMs are powerful for processing and generating text, but they face inherent limitations, particularly in memory and context management. The context window directly impacts how agents interact with users and maintain coherent dialogues over extended periods. This challenge is particularly evident in real-world applications, such as chatbots and embodied agents. Multi-turn conversations with users result in ever-growing interaction histories, which can strain the model's ability to retrieve relevant past information efficiently.

This project aims to leverage Hybrid RAG to enable long-term, persistent, and context-aware interactions between users and agentic systems. By integrating structured knowledge representations with vector-based retrieval, the solution will provide fine-grained control over memory, enhancing both short-term and long-term user-agent interactions.

This project will be conducted remotely and supervised by experts in agentic systems from NVIDIA and the University of Lund.

Competence keywords: Experience of software development in Python

Familiarity with deep learning frameworks

Experience with APIs for LLMs

Connection to WARAs: WARA-Media & Language

Project proposal number: 3

Name of the proposed academic supervisor: Richard Johansson

Name of the proposed industrial supervisor: Samuel Genheden

Project proposal description:
In chemistry, and in drug discovery, several recent publications have shown that it is possible to streamline common tasks, discover new molecules and materials, and integrate with robotics environments by building agents that are powered by large language models (LLMs). Such agents typically augment the reasoning capabilities of the LLMs with external tools and databases. To enable flexibility and scalability, it is advantageous that such systems are multi-agent systems, where specialized agents communicate in some specified workflow to solve a common task.

In this course we will take a practical approach and develop a multi-agent system that can solve tasks in drug discovery. The students will initially work in small sub-groups responsible for building of a specific agent. Such agents could include one capable of generating novel compounds using the Reinvent tool from AstraZeneca, one capable of looking up chemical information in publicly available databases, and one capable of generating Python code to perform simple cheminformatic tasks, to mention a few. When the individual agents are built, the students will explore multi-agent workflows to connect the individual agents and evaluate their capacity, scalability, and emergent properties. We will use local, open-source LLMs that are served from software such as ollama or llama.cpp. Students will be encouraged to take inspiration from existing publications and open-source code in developing their multi-agent system.

Competence keywords: Python programming

Basic ML experience

Connection to WARAs: WARA-Medicine

Project proposal number: 4

Name of the proposed academic supervisor: Michael Felsberg

Name of the proposed industrial supervisor: Leif Haglund

Project proposal description:
Background

Maxar uses its automated multi view stereo algorithms, with direct access to the world's most trusted satellite imagery. This powerful combination creates geospatial capabilities enabling the production of the Globe in 3D. The Globe in 3D provides users with the entire world in highly accurate 3D models for telecommunications, emergency response and defense and intelligence.

The 3D model constitutes of a textured 3D irregular TIN mesh based on ~50cm satellite images,

Project

The last year Maxar has launched 6 new satellites with 30 cm resolution.

The task is to texture a model with these higher resolution images, not only for visualization but also for reasoning like line-of-sight, radio propagation, semantic features (windows, doors …), i.e. Gaussian splatting and Nerfs can be tools but are currently not sufficient.

There are several challenges to achieve this in a global context, e.g.:

• Atmosphere; given images from several days, the atmosphere changes and gives rise to color and contrast differences

• Seasons; is it possible to combine images from different seasons?

• Only partial coverage; in many cases some collection angles are missing – is it possible to infer data from the visible facades?

• Change in model, e.g. new buildings; How to realize that these pixels should not be used when texturing the old ones – and flag that change has occurred.

• Obscurations; Can trees in front of a building be filtered away when texturing the façade?

Competence keywords: General ML and AI competence

Possible AI algorithms to utilize could be generative approaches for (facade) texture generation like, masked auto-encoders, implicit neural representations and/or conditional diffusion models.

Connection to WARAs: WARA-Public Safety

Project proposal number: 5

Name of the proposed academic supervisor: Yiannis Karayiannidis

Name of the proposed industrial supervisor: Matteo Iovino

Project proposal description:
This project focuses on safe dual-arm mobile manipulator navigation performing cart manipulation tasks, in particular handling holonomic utility carts. The project is connected to one of the tasks of WARA Robotics Challenge 2024. The idea is to perform whole-body control addressing both navigation of the robot platform but also steering of the cart to avoid collisions.

Competence keywords: Robotics, SW development, Control and Navigation

Connection to WARAs: WARA-Robotics

Project proposal number: 6

Name of the proposed academic supervisor: George Nikolakopoulos

Name of the proposed industrial supervisor: Matteo Iovino

Project proposal description:
In this project, a group of 3-7 students will develop Reinforcement Learning-based strategies to enable mobile manipulators to safely interact with large complex objects and reconfigure their pose. An example scenario is the reconfiguration of the pose of a wheeled armchair by a mobile manipulator, which is not altogether straight forward even for humans. The student group will use onboard vision, guided by Vision Language Models (VLMs), to a) estimate the current pose of the object in an inertial frame-of-reference, b) navigate close to the target object, while choosing a direction-of-approach suitable for manipulation and c) chose appropriate contact points before deploying a pose-reconfiguration algorithm. The VLMs will also guide high-level decision making to determine if the object is movable or immovable and if reconfiguration of the object is beyond the capacity of the manipulator. The student group will utilize Reinforcement Learning to learn robust policies for mobile manipulators to achieve safe pose-reconfiguration of large complex objects. It is not necessary to strictly follow the above methodology. The student group will be encouraged to think originally about the problem and new ideas will be welcomed.

This project proposal is supported by WARA Robotics.

Competence keywords: This project is looking for a group of students with a diverse set of competencies to collaborate and achieve a challenging objective in the context of indoor robotics. The project seeks competencies in Reinforcement Learning, Computer Vision, Mobile Manipulation, Vision Language Models, Automatic Control, and Manipulation planning. Enthusiasm for experimental work is a soft requirement, as it will help in the execution of the project.

Connection to WARAs: WARA-Robotics

Project proposal number: 7

Name of the proposed academic supervisor: Christophoros Kanellakis

Name of the proposed industrial supervisor: Jesper Tordenlid, Katarina Iversen

Project proposal description:
The development of advanced autonomous robotic systems will lead to a significant enhancement in surveillance operations in maritime environments. The project aims to produce UAVs equipped with state-of-the-art sensors, communication systems, and inspection algorithms. These autonomous systems will demonstrate the ability to perform complex inspection tasks improving the detection and response to maritime activities. The successful implementation of these systems will contribute to the safeguarding of maritime borders, protection of marine resources, and overall enhancement of maritime security.

The research methodology for this project will involve a comprehensive approach that includes theoretical analysis, algorithm development, and experimental validation. The project will commence with an extensive literature review to identify existing methodologies and open-source technologies in the field. Following this, the design and development phase of the mission will focus on creating UAVs with communication systems for robust perception, navigation and control algorithms for surface landing. The inspection tasks will be addressed through the development of real-time image processing, machine learning algorithms, and precise surface landing strategies between the UAVs and USVs.

Challenges to investigate

- Object detection
- Automatic control for surface landing
- Object tracking
- Path planning
- Low bandwidth communication

Competence keywords: Autonomous Robotic Systems, Unmanned Aerial Vehicles (UAVs), automatic control, Image Processing, Machine Learning

Connection to WARAs: WARA-Robotics, WARA-Public Safety

Project proposal number: 8

Name of the proposed academic supervisor: Petter Ögren

Name of the proposed industrial supervisor: Katarina Iversen

Project proposal description:
We, a group of five students from four universities are proposing this project to integrate hardware into our research. As members of WARA-PS, we see this as a great opportunity to utilize its advanced resources. The project explores a simple dynamic shield defense scenario with two drones: Drone A protects a hotspot, while Drone B attempts to breach it. Drone A must physically repel Drone B.

Possible challenges in this project involve developing a robust autonomy framework for the drones, integrating real-time perception (vision-based object detection with fallback to GPS), state estimation (possibly with dynamic factor graph-based SLAM), autonomous decision-making (potentially using Behavior Trees if not MARL), and control strategies for interception. Additionally, scaling from a single-drone defense scenario to multi-drone coordination.

Though student-driven, the project has the support of both academic and industrial supervisors who have agreed to be involved. Additional support consists of the WARA-PS core team and possibly the AILAB at Linköping university. These will lend out the drones for experimentation and demonstration. Moreover, they will assist with integration and evaluation of software into the drones. Since most of the project students conduct research aligning with drone applications, future projects will benefit from reduced overhead and a smoother start from the experience gained in this project course.

Competence keywords: Multi-agent reinforcement learning

Navigation

Autonomous drones

Simulation

Connection to WARAs: WARA-Public Safety

Project proposal number: 9

Name of the proposed academic supervisor: Florian Pokorny

Name of the proposed industrial supervisor: Matteo Ivonio

Project proposal description:
Predicting the effect of physical interactions between a robot and objects in the environment remains one of the grand challenges of robotics and machine learning. Machine learning-based models that utilize training data from past interactions to train a physics model of the interaction are among the most promising approaches in this field due to their ability to capture even hard-to-simulate effects such as friction and interactions with complex environments. The task is to develop a data-driven approach that can predict with high accuracy the pose of an object being pushed along a given target trajectory by a robotic arm. The input to the problem are images taken from a video camera observing the scene as well as the motor commands being sent to the robot. Students will be provided with a large scale dataset of robot/object interactions of randomized pushing interactions collected on the CloudGripper (cloudgripper.org) system at KTH as well as remote access to these robots. Datasets for both rigid object interactions (such as pushing 3d printed objects with various shapes) and deformable object interactions will be made available for testing. As a stretch goal beyond these tasks, students may investigate control approaches utilizing the developed interaction prediction model to complete a pose correction task based on vision sensor data. At the end of the project, students will be able to adapt, deploy and test their algorithms at the Robotics WARA with an ABB YuMi robot.

Competence keywords: This project is intended not just for PhD students interested in robotics but also those from other disciplines. Offline datasets will be provided. If you are interested in foundation models in general, please feel encouraged to join a team working on this project. Multiple separate teams working on this project in parallel are also possible for this project.

Connection to WARAs: WARA-Robotics

Project proposal number: 10

Name of the proposed academic supervisor: Per Runeson

Name of the proposed industrial supervisor: Sigrid Eldh

Project proposal description:
AI tools are introduced in several businesses, including software development. This project aims to explore two aspects of AI tools, namely adoption attitudes and integration of tools in the tool chain. Firstly, we aim to explore how software developers act and react to including AI tools in their workbench. Secondly, we want to investigate how these tools are integrated in the tool chain, both front-end user interfaces and back-end engine integration including the training data for such tools.

Ericsson has piloted AI tools for specific purposes (e.g. https://doi.org/10.1007/s10664-024-10507-y) and is now in the process of rolling out AI tooling for development, requirement handling and maintenance. Developers have been asked to adopt and start using the tools provided during 2025. We want to study the integration and adoption process of AI tooling, what the user experience is, and what, if any, obstacles or hurdles might be present.

We plan to collect data from user interviews and surveys, and, if possible, compare it to tool logs on usage. We plan to present a comprehensive overview of what tools are used for, how the developers experience the tool usage, how they are integrated into the tool chain, and to provide actionable conclusions on how to further leverage AI enhanced tools to increase developers' productivity.

Competence keywords: AI and Machine Learning (6hp)

Software Engineering and Cloud Computing (6hp)

Connection to WARAs: None

Project proposal number: 11

Name of the proposed academic supervisor: Björn Olofsson

Name of the proposed industrial supervisor: Johan Silvander

Project proposal description:
The radar is the most important sensor when it comes to safe navigation in a maritime environment.

However, since the radar is primarily used by human operators the tuning of the radar parameters are performed by humans with the help of heuristics.

Since we are operating in an autonomous unmanned environment, with no human in the loop, the tuning of the radar parameters has to be performed with the help of some sort of optimization algorithms.

The tuning is dependent on, for example weather conditions, sea state, type of task that shall be performed etc.

The algorithm shall be implemented in a way that it is possible to executed on an edge device with the performance less/equal of an Nvidia Jetson Orin Nano.

Competence keywords: Optimization, AI, autonomy, software engineering, robotics, unmanned

Connection to WARAs: WARA-Public Safety

Project proposal number: 12

Name of the proposed academic supervisor: Johanna Björklund

Name of the proposed industrial supervisor: Konrad Tollmar

Project proposal description:
The main task of this project is to define an evaluation framework for Game LLMs. You will use platforms and tools like Unity and/or Godot to build and test evaluation methods using available open-source games and LLM models.

Large language models (LLMs) are revolutionizing video game development, offering the potential to create dynamic and engaging experiences unlike anything seen before. These powerful AI models can generate realistic dialogue, craft intricate narratives, and even control agents in gameplay. However, as the applications multiply, so does the importance and complexity of measuring performance.

Evaluating LLMs in Games is an evolving field, as these systems go beyond simple text evaluation to also include assessment of actions and decisions, i.e., the performance of LLM in gameplay. This requires new metrics and methods. Here are some examples:

- Dialogue Consistency in NPC Interaction

- Simple Quest Generation with Clear Goals

- Basic Action Appropriateness in a Controlled Environment

- Helpfulness of a hint system

These examples demonstrate a shift from purely linguistic evaluation to assessing the LLM's impact on the overall player experience. One approach is to use Human Evaluation to get how Players understand and rate performance. This is similar to how LLM are trained with Reinforcement Learning from Human Feedback (RLHF).

Competence keywords: You will use platforms and tools like Unity and/or Godot and build tests using open-source games and LLM models. Moreover, good knowledge of Python and familiarity with LLM models and tools, as well as general knowledge of Human-Computer interaction, are needed.

Connection to WARAs: WARA-Public Safety

Project proposal number: 13

Name of the proposed academic supervisor: Michael Felsberg

Name of the proposed industrial supervisor: Fredrik Viksten

Project proposal description:
This PhD-level project centers on advanced techniques for 3D-reconstructing the interiors of historic buildings, prioritizing sparse viewpoints while achieving robust, high-fidelity results. This course focuses on pipeline design, data strategies, and the scientific evaluation of methods, not the specific end result but it should still be possible to show case a resulting model from at least one site.

Objectives:

1) Devise a pipeline for accurate indoor 3D modeling from limited viewpoints, ensuring resilience to occlusions and uneven lighting.

2) Develop best practices for data collection (minimal vantage points, controlled angles, lighting) to reduce manual effort and maximize reconstruction performance for your developed pipeline.

3) Optionally incorporate object detection and segmentation using 2D texture, 3D mesh analysis, or hybrid methods to isolate individual artefacts or structural features.

4) As an optional continuation of point 3 above, enable seamless updates of textures with recolored or deoldified imagery, facilitating iterative enhancements to enable heritage reconstruction.

Data from multiple heritage sites will be available for testing, but students may need to gather supplemental images to validate their techniques. The final outcome is a robust, well-documented system that furthers cultural heritage preservation and pushes the boundaries of 3D computer vision research.

Competence keywords: Sparse-View 3D Reconstruction

Structure-from-Motion (SfM)

Neural Radiance Fields (NeRF)

Depth Estimation & Feature Matching

Photogrammetry & Multi-View Geometry

Gaussian Splats

3D Scene Understanding

AI-Based Image Restoration

Texture Transfer & Recolorization

Mesh Processing & 3D Visualization

Connection to WARAs: WARA-Public Safety

Project proposal number: 14

Name of the proposed academic supervisor: Mariusz Wzorek

Name of the proposed industrial supervisor: Katarina Iversen

Project proposal description:
In this project, we aim to integrate concepts from human-aware robot navigation into robotic exploration tasks in the context of search and rescue (SAR) operations. The developed methods should enable robots to adapt their exploration strategies by explicitly considering human actors' (such as rescue operators and potential victims) presence, behavior, and intentions. This adaptation is expected to ensure that exploration tasks are carried out safely and efficiently while enabling effective collaboration with the humans involved in the SAR domain (i.e., improving cooperation with first responders or increasing the likelihood of successful victim detection and assistance).

Example research topics targeted in the project include vision-based human detection, tracking and activity recognition, multi-modal sensor fusion, temporal motion planning, trajectory prediction models of dynamic objects, SLAM with dynamic obstacles, and realistic human motion models. The starting point will be an existing 3D exploration method for large-scale environments available as a software module.

The project will be performed within the WARA-PS arena with the supervision of researchers from the AIICS division at Linköping University. Project results in the form of software modules are expected to be deployed and demonstrated in simulations, and in the field robotic experiments using one of the ground/air robotic systems, including Clearpath Husky/TurtleBot4, Boston Dynamic Spot, or DJI M 100/300.

Competence keywords: vision-based object detection and tracking

motion planning

sensor fusion

SLAM

Connection to WARAs: WARA-Public Safety