Projektvalskatalog

EF112X kandidatexamensarbete inom elektroteknik (15 hp) våren 2027

I denna katalog kan du hitta information om alla valbara projekt inom kursen EF112X kandidat-examensarbete inom elektroteknik (15 hp) som erbjuds våren 2027 vid EECS skolan, KTH. Kursen EF112X pågår hela vårterminen från mitten av januari till slutet av maj. Projekten utförs i grupper om två. I år erbjuds 72 projekt att välja emellan inom ramen av 17 olika kontext.

Kontext inom systemteknik och robotik

- Kontext A: The dynamics of a sustainable society (*Matthieu Barreau*)
- Kontext B: AI-driven Power Grids (*Giuseppe Belgioioso*)
- Kontext C: Information & algorithms in social network dynamics (*Angela Fontan*)

Kontext inom inbyggda system och elkraftteknik

- Kontext D: Electric transportation (*Nicholas Honeth*)
- Kontext E: Power systems control (*Qianwen Xu*)
- Kontext F: Hydro power modelling (*Mikael Amelin)*
- Kontext G: HVDC supergrids for offshore wind (*Ilka Jahn*)

Kontext inom elektromagnetism, fusion och rymdteknik

- Kontext H: Novel microwave/antenna technologies (*Oscar Quevedo Teruel*)
- Kontext I: Fusion solens energikälla på jorden (*Thomas Jonsson*)
- Kontext J: Solar wind and planetary environments (*Lorenz Roth*)

Kontext inom information och nätverksteknik

- Kontext K: Networked Artificial Intelligence (*Carlo Fischione*)
- Kontext L: Cybersecurity (*Hamed Nemati*)
- Kontext M: Information engineering: Big Data & AI (*James Gross*)

Kontext inom datavetenskap och maskininlärning

- Kontext N: AI, games, and strategy (*Mika Cohen)*
- Kontext O: Computational brain modelling (*Pawel Herman*)
- Kontext P: Sign language processing (*Jonas Beskow*)
- Kontext Q: Software Integrity (*Martin Monperrus*)

Tillgängliga projekt 2027

Alla projekt som erbjuds vårterminen 2027 finns beskrivna i denna pdf-fil. Läs igenom projektbeskrivningarna noggrant. I valet markerar du vilka projekt du helst vill göra. Obs, inom ramen av denna kurs kan man inte "skräddarsy" sitt eget projekt eller göra ett industriprojekt. Du måste välja ett av de tillgängliga projekten som finns beskrivna i denna katalog.

Viktiga datum

- **Informationsmöte om projektvalet:** onsdag, 7 okt 2026, kl 13:00-17:00, sal D2
- **Projektval:** 1-15 nov 2026
- **Kursstart:** vecka 3, 2027 (uppstartmöte, i anslutning första träff med Supervisorn)
- **Kursslut:** vecka 21, 2027 (heldag KEX-presentation)

Kurs-PM
All information om kursens uppbyggnad finns i kurs-PM. Kurs-PM:et läggs ut senast den 1 november 2026 på projektvalsidan: www.kth.se/social/course/EF112X/

Behörighet
För studenter som läser elektroprogrammet gäller: minst 104 högskolepoäng från kurser i utbildningsplanen, till och med period 1 i årskurs 3, ska vara avklarade senast vid startdatum för period 2 för att studenten ska få påbörja kandidat-examensarbetet. Om du kommer från en annan KTH skola gäller denna skolans regler om behörighet.

Anmälan till KEX-kursen
Elektroteknikstudenter: Om du går i årskurs tre (CELTE-3) och ligger i fas med studierna, ska du välja villkorligt valfria, och/eller den helt valfria kursen, och kandidatexjobbskursen EF112X mellan 1-15 november via antagning.se. Logga in med ditt kth.se konto.

Om du antogs till elektroteknik 300 hp år 2021 eller tidigare, sker anmälan till kandidat-exjobbskursen EF112X istället via e-post till svl-celte@kth.se mellan 1-15 november.

Om du läser ett annat kandidatprogram (fysik, farkost, teknisk matematik, data eller energi & miljö), anmäler du att du vill göra kandidatexjobbskursen EF112X till studievägledaren vid respektive KTH skola.

Val av projekt
Förutom att anmäla dig till kandidatexjobbskursen behöver du (oberoende från vilken skola du kommer) välja på internet vilken av kandidatexjobb-projekten du helst vill jobba med.

När sker valet?
Valet av kandidatexjobbprojekten görs under perioden 1-15 november 2026. Resultatet påverkas ej av när du väljer under valperioden.

Projektgruppen
Kandidatexjobbsprojektet utförs i grupper om två studenter. Om du inte lyckas hitta en projektpartner, tilldelas du en partner med liknande projektönskemål som du.

Gör ditt val
Anmälan görs på projektvalsidan (www.kth.se/student/kurser/kurs/EF112X). Välj de sju projekten ur denna projektvalskatalog som intresserar dig mest. Du kommer kunna ange din prioriteringsordning när du väljer (prio 1= projektet du helst vill ha). Om du redan har hittat en projektpartner, fyll i bådas namn, e-mail och program i samma anmälan. Gör endast **en** anmälan per grupp. Om du inte har en projektpartner än, anmäl dig ensam (du kommer sedan tilldelas en projektpartner).

Lycka till med projektvalet!

Anita Kullen (kullen@kth.se)
Kursledare för kandidatexjobbskursen EF112X

Stockholm, 5 oktober 2026

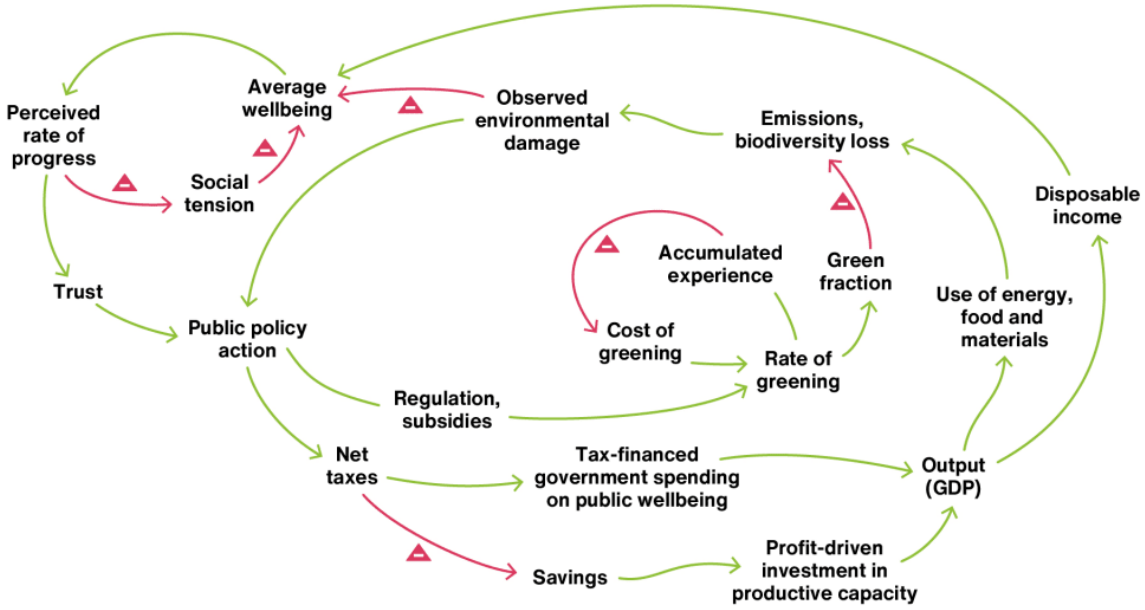
Context A:
The Dynamics of a Sustainable Society

Context Responsible: Matthieu Barreau, barreau@kth.se Decision and Control Systems

"*Meeting the needs of the present without compromising the ability of future generations to meet their own needs.*"

This is the well-known sustainability definition by the United Nations Brundtland Commission in 1987. Sustainable development takes into consideration environmental concerns along with economic development and social aspects. The 17 Sustainable Development Goals are a framework for improving the lives of populations around the world and mitigating the hazardous man-made effects.

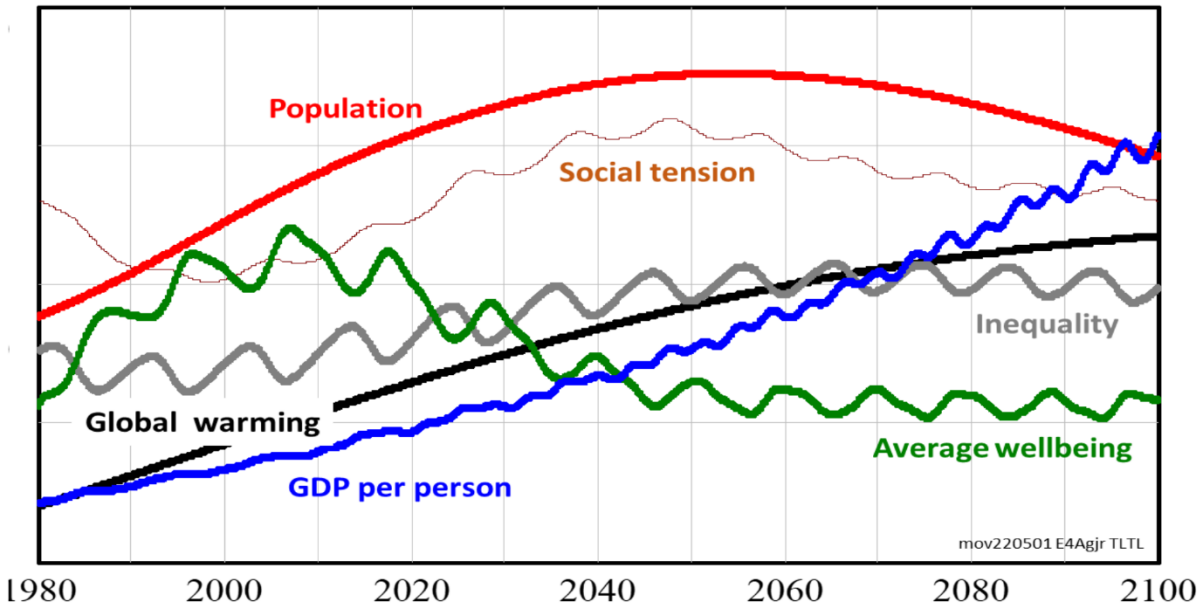
But how is this done? We are dealing with an extremely complex and complicated problem that involves the law of nature, global politics and economics, cultures, human behavior, and technical solutions, to name just a few of the aspects. And we need to deal with long time scales and complex interplay between different domains. In these projects, we will approach this problem by dynamical systems modelling.



The World3 model, presented in the book "The Limits to Growth", was a pioneering computer-based system dynamics model that aimed to shed light on the complex interactions between population growth, resource depletion, industrialization, and environmental impacts. The central message of the model was that if left unchecked, human activities such as rapid population growth and excessive resource consumption would eventually lead to ecological and societal challenges. This model has been updated multiple times and totally rewritten in 2022 with the new name Earth4All.
It generated heated debates and criticism, particularly regarding the assumptions made, prediction capabilities, and policy implications. While the model had its limitations, it played a pivotal role in raising awareness about the challenges of sustainable development and the importance of considering the finite nature of Earth's resources in shaping our future. The ongoing debates surrounding the model continue to shape discussions on environmental and economic policy.

Projects A1 – A3: Modeling the world

This context has three projects. The aim of these three projects is to analyze and experiment with the Earth4All model. The focus of each project is different, but there is a common ground. There will be three independent projects, A1 – A3, but you will interact with the groups in multiple workshops, where you will share knowledge and identify connections between the methods.



The Earth4All model is a dynamic model with interconnected subsystems describing: 1) food production and agriculture, 2) industrial production, 3) human population, 4) non-renewable resources, and 5) pollution. A typical simulation of the model is shown at the top. The initial step will be common for all projects. The students will first make use of a web graphical interface of the Earth4All model to understand its structure and build some intuitive understanding. In the later phase, you will use and modify the Python implementation of the model.

Project A1: Modeling the world – Exploring possible futures with systems thinking

Supervisor: Matthieu Barreau, barreau@kth.se, Division of Decision and Control Systems

In this project, you will work with the Earth4All model to explore possible futures of our world. You will work on a subpart of the model since it contains many variables that are dynamically interconnected in reinforcing and stabilizing loops. You will use systems thinking and control/systems theory to analyze it and to propose and evaluate aspects with an implication on sustainability.

(1) Analyze the causal loop diagram to understand its structure. Find the main relations between the variables of the model. Identify important loops. Define the variables and the external inputs.
(2) Use a mix of simulation and analysis to understand the behavior of the model. Use sensitivity analysis to identify variables and inputs of importance. Identify important delays and inertia that affect the dynamic response of the model.
(3) Based on your systems understanding, explore different possible futures, for example, new policies, changed behavior, technology development, or other important factors. Use simulation and analysis to explain the effects. Experiment with the model and modify it, if necessary, by changing parameters or introducing/removing links.

Project A2: Modeling the world – Model identification using physics-informed learning

Supervisor: Matthieu Barreau, barreau@kth.se, Division of Decision and Control Systems

The dynamics of one sector are very complex and not written in a form suitable for control. The tool you will use is Physics Informed Learning.

1) You will first assume no knowledge of the original system and try to find a nonlinear dynamical model using a neural network. You will need to collect data and train the model.
2) The second step is to try to simplify the neural network from the previous part. You can investigate the dimension of the original system, but also try to simplify the model by identifying key signals.
3) You will go deeper into the identification by considering submodels and repeating the first step with this new system design.
4) From your understanding of the different variables, you will try to find which ones are measurable. The final identified model should be interpretable and help you to highlight the most important features and the policies with the most impact.

Project A3: Modeling the world – Optimal policies with reinforcement learning

Supervisor: Matthieu Barreau, barreau@kth.se, Division of Decision and Control Systems

Reinforcement Learning (RL) is a well-known machine learning technique often used to find an optimal policy for controlling a dynamic system in order to maximize a reward cumulated over time. At each time step, the agent takes an action, influencing the system's evolution over time. The agent then observes the updated states and receives a reward based on the transition. RL is often used in uncertain environments where the dynamics of the system need to be learned while interacting with the system itself. Over time, the agent gathers more information, enhancing its understanding of the system dynamics for more informed decisions. RL is used for many applications such as robotics, automation, video games, and finance and recommendation systems. A similar tool that can be used is Model Predictive Control (MPC)

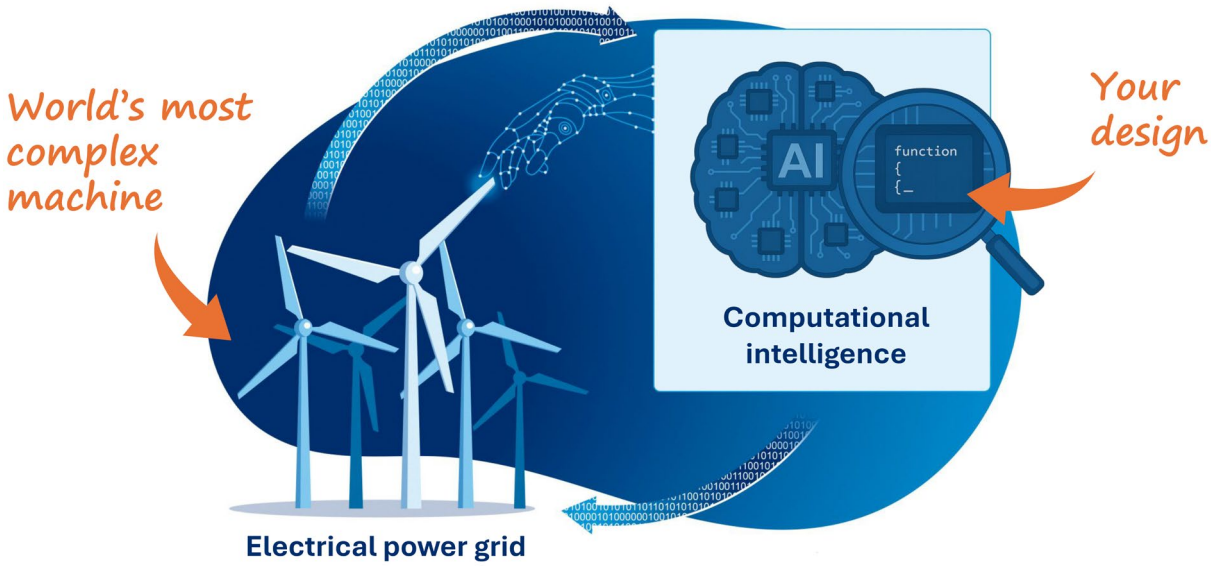
In this thesis, we propose the application of RL techniques to a subsystem of the World3 model, a well-known system dynamics model used for studying global sustainability. Using RL, the students should try to derive the optimal control policy with respect to a given cost. Constraints on the measurements (availability of the measure, sampling…) and the control inputs (delayed inputs, quantization…) can be considered. More precisely, the aim is to optimize resource allocation and policy decisions within the selected subsystem, with the overarching goal of contributing to a specific UN sustainability goal. The World3 model provides a comprehensive framework for simulating interactions between population, resources, and the environment, making it a good testbed for the students' RL experiments.

This thesis will require a thorough understanding of the structure and dynamics of a subsystem within the Earth4All model.

1) Identify relevant state variables, actions, and their interdependencies.
2) Formulate a reward function that can well quantify the desired system behavior linked to a specific UN sustainability goal.
3) Train and evaluate the RL agent using appropriate algorithms, in order to derive an optimal control policy.
4) Assess the performance of the obtained control policy with simulation experiments and compare it against a baseline approach.

Context B: AI-driven Power Grids

Context Responsible: Giuseppe Belgioioso (giubel@kth.se)
Division of Decision and Control Systems



The Challenge

The electrical power grid - world's most complex machine - is under unprecedented pressure. To meet climate goals and strengthen energy independence, massive amounts of renewable generation are being connected. Unlike traditional power plants, however, renewables are dispersed at the edge of the grid, and their output rises and falls with the weather.

This variability collides with rapidly growing electricity demand: the surge of energy-intensive data centers, the electrification of transport, and record-breaking heatwaves driving air conditioning loads. The result is a grid operating at its limits. Recent large-scale blackouts, including the world-record outage across the Iberian Peninsula in April 2025, have revealed how fragile the system can be. When bottlenecks occur, clean energy is wasted through curtailment, while overloaded lines and over-voltages threaten the grid's safe and reliable operation.



Figure 1: Automation will play a key role in shaping future sustainable energy systems.

At stake is nothing less than the backbone of the energy transition. The central question is no longer whether renewables can be integrated, but how can we fundamentally redesign grid operations to ensure security of supply in a world powered by variable resources.

Industry state-of-the-art manual and semi-automated mechanisms for congestion control in power transmission grids are unsuited for these new tasks for two main reasons:

1. They are not fast enough to safely respond to generation variability.
2. They cannot handle the complexity of large number of small-size generators.

Consequently, there is a growing need to enhance real-time automation and implement control actions at shorter intervals, ideally every few seconds. The Swedish system operator Svenska Kraftnät estimates potential savings in the order of tens of billions of SEK over a decade if real-time control of power flows can prevent the construction or reinforcement of power lines. For the same reason, in the United States, the connection of more than 2000 GW of renewable generation is currently being delayed due to grid capacity constraints.

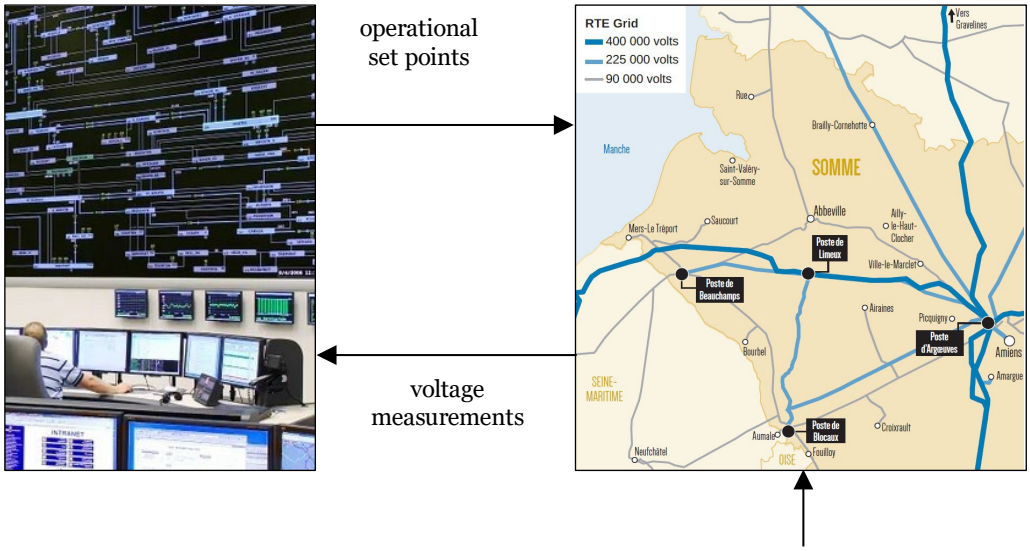


Figure 2: In future electrical power grids, generators and other grid assets will be autonomously controlled in real-time via smart algorithms (i.e., AIs) and without any human intervention.

Your Mission

You will explore how AI, namely, smart optimization and control algorithms, can help manage congestion in the world's largest and most complex machine: the electrical power grid.

By the end, you will have learned how to:

- Understand the basics of power grids including the power flow equations, governing the electrical power grid dynamics, and the Pandapower (link), an open source Python tool to simulate how electricity moves through the grid.
- Use optimization and control principles essential for designing smart congestion-control algorithms, and gain hands-on experience with state-of-the-art optimization solvers such as IPOPT (link);
- Test your own ideas in simulations using real grid models, from local to continental size, and real data to validate how well your algorithms work in practice.

Remarks:

- Prior knowledge of power systems and optimization methods is not required. The scope of the project can be adapted online to match background and interests.
- Projects delivering promising results can be potentially turned into publications.

Project B1: Closing the Loop: Feedback Optimization of Grids

Supervisor: Giuseppe Belgioioso (giubel@kth.se), Division of Decision and Control Systems

Industry state-of-the-art congestion control mechanisms heavily rely on grid models and forecasted disturbances to manage power flow. In control theory, this is referred to as feedforward control. The challenge arises when grid models or forecasts are inaccurate, leading to deviations between expected and actual outcomes. In practice, grid operators need to update forecasts and take corrective actions, a process known as re-dispatching.In this project, you will:

1. Familiarize yourself with the simulation environment;
2. Develop a state-of-the-art feedforward controller for congestion management;
3. Validate the controller design in the simulation environment, exploring the impact of forecasts errors on the quality of the outcomes;
4. Explore alternative feedback control designs that mitigate these issues.

Project B2: Scaling Up: From Local to Continental Grids

Supervisor: Giuseppe Belgioioso (giubel@kth.se), Division of Decision and Control Systems

Industry state-of-the-art congestion control mechanisms are based on mathematical optimization. Specifically, power generation set-points for the upcoming day are determined by solving an nonlinear program (NLP) of the following form

$$\begin{aligned} \min_{s} \quad & \text{curtailment}(y) \\ \text{s.t} \quad & y = h(s,d) \\ & y \in \text{limits}, \end{aligned}$$

where *s* represents the generators set-points, *y* the grid's state, *d* the disturbance, and *h* is the grid model, i.e., the outcome of the power-flow equations. Notably, the dimension of this optimization problem, depends both on the number of decision variable *s* and on the grid size.

In this project, you will:

1. Familiarize yourself with the simulation environment;
2. Develop a state-of-the-art congestion control algorithm;
3. Validate your design using the simulation environment;
4. Numerically explore the computational complexity of your smart control mechanism by deploying it on grids of increasing size, from local grids to continental grids.

Project B3: Price of Anarchy: Centralized vs Multi-area Grids

Supervisor: Giuseppe Belgioioso (giubel@kth.se), Division of Decision and Control Systems

Industry state-of-the-art congestion control mechanisms treat the transmission grid as an isolated system. However, modern transmission grids consist of multiple interconnected areas, each locally managed by different operators but physically interconnected. A natural question in this more realistic scenario is whether the multi-area grid, resulting from the interconnection of locally controlled areas, can retain efficiency. In this project, you will:

1. Familiarize yourself with the simulation environment;
2. Develop a state-of-the-art congestion control algorithm;
3. Validate your design using the simulation environment on a single-area grid;
4. Quantitatively compare the loss of efficiency (i.e., the price of anarchy) and reliability of your control algorithm in a centralized vs multi-area scenario.

Context C: (Mis)information and Algorithms in Social Network Dynamics

Context Responsible: Angela Fontan, angfon@kth.se, Dept. Decision and Control Systems

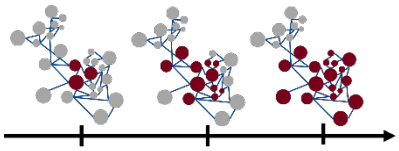
Social networks play a central role in shaping opinions and behaviors in society. They are often modeled as graphs, where individuals are represented as nodes and their relationships as edges, and, over the past decades, social network dynamical models, including models of opinion formation, have been developed to capture how the states (e.g., opinions and beliefs) of interconnected individuals evolve over time through social interactions.

Studying these networks and understanding the spread of information (or misinformation!) over them poses significant challenges. Questions of interest are, for instance: What risks arise when malicious agents attempt to destabilize communities by spreading misinformation? Are modern recommender systems and algorithms contributing to the spread of misinformation and exacerbating polarization? What mitigation and intervention strategies can be designed to address these challenges?

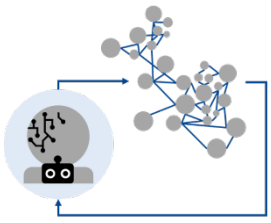


The three projects in this context address these questions from complementary perspectives. They investigate the dynamics of social networks, the spread and mitigation of misinformation, and the role of recommender systems in shaping online communities. The projects make use of both synthetic networks and real-world datasets to develop and analyze models through numerical simulations and to illustrate and evaluate the resulting findings.

Projects C1 – C3



This context consists of three projects, C1 - C3. These three projects aim to analyze and experiment with synthetic and real-world datasets on social networks. They are focused on different aspects of exploring complex challenges in social network dynamics and recommender systems. However, you will all share common datasets and the same dynamical model for opinion evolution; moreover, you will all discuss potential methodologies and the social implications of your findings. Suggested literature will be provided by the context responsible[1].



[1] Examples include:

- N. Pagan, A. Philippou, G. De Pasquale, “Misinformation Mitigation over Social Networks: A Control Approach”, Proc. 8th Annual Learning for Dynamics and Control Conference, PMLR 331:1954-1965, 2026
- G. De Pasquale, S. Dean, P. Frasca, “Recommender Systems as Control Systems”, arXiv:2605.01503, 2026

Project C1: Engagement and Misinformation Trade-offs in Social Networks

Supervisor: Angela Fontan, angfon@kth.se, Department of Decision and Control Systems

Project C1 investigates the relationship between user engagement and misinformation in social networks. The project focuses on understanding how network structure and the presence of radical users affect the effectiveness and consequences of misinformation mitigation strategies.

1. Study measures of misinformation propagation and user engagement in opinion-dynamics models;
2. Investigate how different network structures and distributions of user opinions influence the spread of misinformation and the effectiveness of interventions;
3. Analyze the trade-offs between reducing misinformation and maintaining user engagement, including the effects of interventions on different groups of users;
4. Explore these phenomena through simulations using synthetic networks and real-world datasets.

Project C2: Misinformation Mitigation through Social Network Control

Supervisor: Angela Fontan, angfon@kth.se, Department of Decision and Control Systems

Project C2 investigates how misinformation can propagate through social networks and how interventions can be designed to reduce its spread while preserving user engagement. A controlled opinion-dynamics model is used to study the interaction between misinformation, engagement, and network structure.

1. Study the spread of misinformation in social networks and the role of content characteristics (e.g., sentiment, novelty);
2. Study the extension of well-known opinion dynamics models to incorporate misinformation mitigation and engagement-aware control;
3. Implement model-based vs. model-free control strategies and investigate their behavior under different network structures and user configurations;
4. Validate the methods through simulations using synthetic networks and datasets.

Project C3: Fairness and Feedback in Recommender Systems

Supervisor: Angela Fontan, angfon@kth.se, Department of Decision and Control Systems

Project C3 investigates how recommender systems can create feedback loops that amplify polarization, reduce content diversity, and produce unequal exposure for content creators. A control-theoretic model of the interaction between users, recommender systems, and creators is used to study how different recommendation policies affect the long-term behavior of the system.

1. Study feedback loops in recommender systems and their effects on users and content creators, including opinion polarization, content diversity, and popularity bias;
2. Study how fairness can be incorporated into recommendation policies;
3. Investigate the trade-offs between short-term engagement and long-term outcomes,
4. Using dynamical simulations, study how different recommendation policies affect polarization.

Context D: Electric Transportation

Ansvarig för kontexteten: Nicholas Honeth samt vice-ansvarig Mats Leksell
Institutionen för elkraftteknik

Introduktion

Dagens samhälle är starkt beroende av ett fungerande transportsystem som ska klara av såväl människor som gods. Dock står transporterna på global nivå för en stor andel av världens koldioxidutsläpp i och med att det är framförallt olja som används som bränsle. Den generella lösningen för att minska koldioxidutsläppen heter idag *elektrifiering*. Genom att elektrifiera kan vi både få bättre prestanda och sänka koldioxidutsläppen.

I kontext D studeras en hyperavancerad racingbil samt eldrivna fritidsbåtar och i centrum av kontextet finns batteriet. Ett av projekten undersöker hur en säker laddinfrastruktur kan skapas i marin miljö. Två av projekten studerar själva designen av batterisystemet för en racingbil.

Tillsammans ska projektdeltagarna även sätta in sina system i det överordnade transportsystemet och reflektera över deras betydelse.



Project D1: Electrical Safety and Isolation in Marine Charging Stations

Supervisors: Nicholas Honeth (honeth@kth.se), Mats Leksell (leksell@kth.se)

Background

This bachelor's thesis will contribute to an ongoing research project which aims to address a collection of challenging questions regarding the electrification of light marine transportation on Sweden's coast and inland waterways. Light watercraft such as boats in the order of 6-12m are used extensively for personal transport, service delivery such as medical and police, transportation of goods and building material for a vibrant archipelago, a *levande skärgård*, and efficient use of inland waterways.



Given the high voltages and large currents found in EV charging systems it is essential that electrical safety aspects are rigorously investigated and handled from the marine perspective. Risk factors such as conductive salt water in contact with aluminium hulls create possible scenarios which are not encountered in terrestrial automotive applications.

Objectives

This bachelor's thesis project aims to assist the research project work by investigating scenarios where the use of standard road vehicle charging systems become problematic. Examples of this include:

- The existence of galvanic circuits between boat hull, water and electrical grounding on the dock.
- Performance of earth leakage protection and insulation monitoring devices in marine contexts.
- Interaction of onboard electrical protection devices with the protection systems of the fixed installations on land.

The outcomes should include a good overview of the safety and protection systems used in electric vehicle charging as applied to marine environments. An existing laboratory mock-up is to be extended in order to demonstrate one or more of the issues as a way to communicate the research project results.

Project D2: Efficiency Measurements for Marine Applications

Supervisors: Nicholas Honeth (honeth@kth.se), Mats Leksell (leksell@kth.se)



Background

This bachelor's thesis will contribute to an ongoing research project which aims to address the measurement problem when it comes to electric drives in a marine environment. To measure electric power is a standard procedure but to get efficiency the mechanical power is also needed and that is more problematic in a marine environment. Given the inherent high energy consumption for boats means that an even so small increase of the efficiency will have a big impact on the cruising range.

KTH's Electric Power and Energy Systems lab has an underwater dynamometer for testing submersible electric motors used in boats. A motor sits in a water tank and drives a load machine through a torque transducer. Previous projects have tested propellers, tuned motor drives with machine learning and analysed waveforms on this equipment. In this project you will further analyse how efficiency can be obtained in a marine environment and build a data platform for automated testing.

The platform will support tests with Zparq, a Swedish developer of submersible motors, and a parallel master's thesis on automatic motor-drive tuning. You will work with real hardware, combining measurement technology, electric machines, software engineering and data visualisation.

Objectives

This bachelor's thesis project aims to assist the research project work by designing an automated test bench. The work will include both hardware and software.

Research questions may include:

- How can propulsive power be measured for a leisure craft in a cost-efficient way?
- How is efficiency visualized for a marine environment?
-

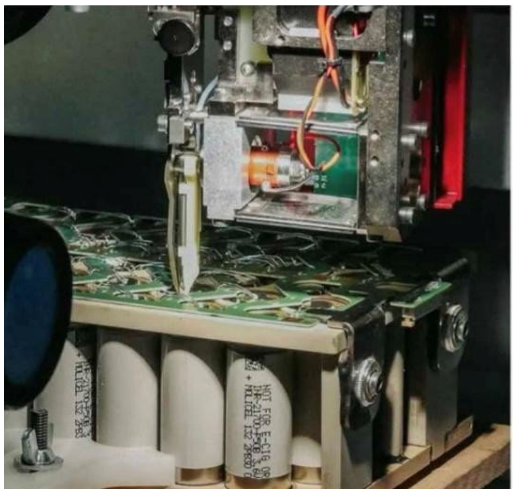
The outcomes should include a good overview of the measurement challenge in marine environments. An existing laboratory test setup is to be extended in order to demonstrate the possibilities for an automated system.

Project D3: Design of a High Voltage Battery Segment using Cylindrical Cells

Supervisor: Carl-Mikael Zetterling (bellman@kth.se) together with the powertrain and electronics team at KTH Formula Student.

Background

Formula Student is an international engineering design competition where university students design, build and race formula-style cars. The KTH Formula Student car is an electric race car featuring a custom-built high voltage battery. Like many other Formula Student teams, the KTHFS team currently uses pouch cells to construct its battery segments. While pouch cells are currently permitted by the competition rules, safety concerns may lead to future restrictions on their use. Such a rule change would require the team to redesign the HV battery using alternative cell types.



KTH FORMULA STUDENT

The Project

This project focuses on the design of a high voltage battery segment using cylindrical cells. The segment should be designed as part of a complete battery pack, considering the required voltage and capacity, packaging constraints, and the integration and interconnection of multiple segments within the battery compartment.

Research Questions

- How should a battery segment using cylindrical cells be designed to meet the electrical requirements of the accumulator?
- How can the segment be designed and integrated to achieve an efficient and practical battery pack within the available constraints?

Tasks

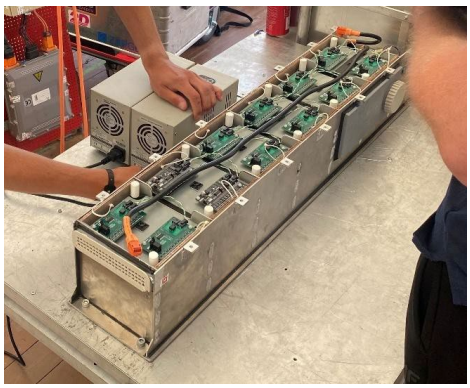
- Research and selection of suitable cylindrical cells and segment configurations in compliance with FS rules.
- Design a complete battery segment, including the mechanical structure and electrical interconnections.
- Construction and evaluation of a prototype battery segment.

Project D4: Design of a Cell balancing Station for a High Voltage Battery Segment

Supervisor: Carl-Mikael Zetterling (bellman@kth.se) together with the powertrain and electronics team at KTH Formula Student.

Background

Formula Student is an international engineering design competition where university students design, build and race formula-style cars. The KTH Formula Student car is an electric race car featuring a custom-built high voltage battery. As Formula Student cars become more performance-oriented, efficient and reliable cell balancing becomes increasingly important. Proper cell balancing maximizes the HV battery's capacity, performance and reliability.



The Project

The high-voltage battery system of the KTH Formula Student race car consists of multiple sub-segments. The scope of this project is to design a dedicated cell-balancing PCB capable of connecting to an individual segment to redistribute energy from higher-voltage cells to lower-voltage cells. This active balancing mechanism ensures equalized cell voltages across the battery segments, thereby maximizing the total usable capacity and overall pack performance.

Research Questions

- How can an active balancing circuit be designed to handle the current cells used within the KTH Formula Student race car?
- How can the cell balancing station be packaged to easily connect to a segment from the HV battery?

Tasks

- Schematic design of the cell-balancing circuit and selection of components optimized for scalability and batch manufacturing.
- Design the PCB layout adhering manufacturability principles to ensure easy replication for future segment layouts.
- Assemble a PCB with the cell balancing circuit integrated, test its performance using a HV segment from DeV20 and document the results.

Context E: Power Systems Control

På svenska

Som ett svar på utmaningen med klimatförändringarna elektrifieras allt fler delar av samhället, och elproduktionen ställs om till att inkludera stora mängder förnybar och delvis distribuerad kraftproduktion. Dessa nya kraftkällor t.ex. vind och solkraft ersätter de större centrala produktionsenheter som tidigare utgjort ryggraden i systemet. Detta ställer nya krav på elkraftsystemet dels eftersom vindkraft och solkraft inte är i lika hög grad styrbart och därför kan påverka systemets stabilitet på nytt sätt. Dessutom är de förnybara kraftkällorna inte lokaliserade i närheten av stora last-centra, exvis. städer, vilket gör att kraven på överföring av elenergi förändras. De senaste åren har dessa frågor blivit alltmer aktualiserat i samhällsdebatten

De nya kraven finns både på transmissions och distributionsnivå. Det inkluderar nya gränser för stabilitet i systemet på grund av minskad roterande massa i generatorer och ökade variationer gällande spänning, effektflöden och frekvens. Dessa nya krav möts effektivast med nya kontroll och automationssystem och även nya styrbara kraftsystemkomponenter, vilka blir allt viktigare för ett välfungerande elkraftsystem. För att dessa kontrollsystem ska fungera krävs mer omfattande mätning och insamling mätvärden från större delar av systemet.

Detta kontext behandlar nya metoder och tekniker för styrning av elkraftsystem med stora mängder förnybar kraftproduktion. Projekten i kontexten inkluderar både traditionella elkrafttekniska frågor såväl som utmaningar inom automation och reglerteknik samt de informations och kommunikationssystem som är nödvändiga för denna automation.

In English

For several reasons, the power system is currently developing to include large amounts of renewable and distributed generation that in part replaces the large central production units that previously formed the backbone of the system. These new distributed power sources, place new demands on the power system partly because they are not equally controllable - but also because they are not always located in the vicinity of large loads.

These changes place new demands on the power system, both at transmission and distribution level. These new demands include new limits on the stability of the system due to reduced rotating mass of generators and increased variations of voltage, power flow and frequency. These new requirements are in turn met most efficiently with new control and automation systems and new controllable power system components, which are becoming increasingly important for an efficient power systems.

This context deals with new methods and techniques for the control of power systems with large amounts of renewable power generation. The projects in the context includes both the traditional electric power issues as well as challenges in automation and control technology and information and communication systems necessary for this automation.

Project E1: Voltage Control of Renewable Energy Integrated Power System

Supervisor: Qianwen Xu, qianwenx@kth.se, Electric Power and Energy Systems

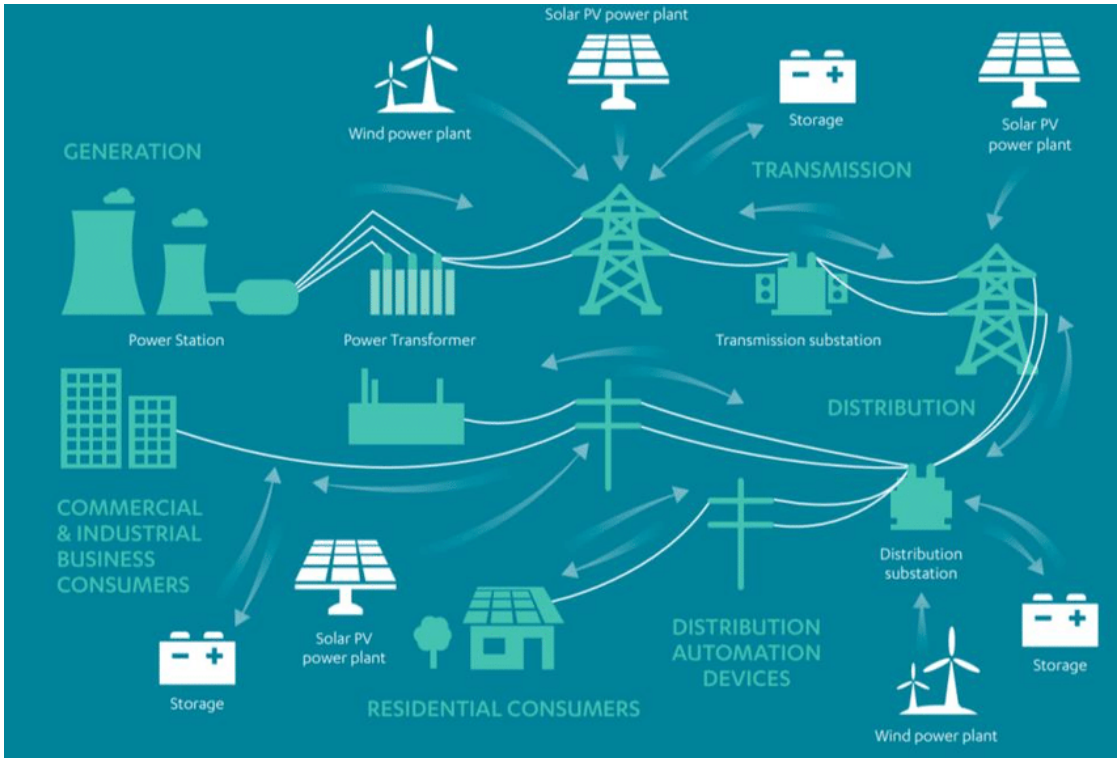


Figure: Future power system

As the penetration of renewable energy sources (RES), such as solar photovoltaics (PV) and wind power, continues to increase in modern power systems, the control of system voltage has become a critical challenge. Traditional voltage control methods cannot handle the variability and intermittency of renewable energy. Voltage fluctuations can lead to power quality issues, equipment malfunctions, and even system instability. To address these challenges, modern power systems must implement advanced control strategies that integrate the reactive power control capabilities of inverter interfaced renewable energy sources, along with other flexible resources such as battery energy storage, electric vehicles, demand response assets.

This project will investigate the impact of high levels of renewable energy integration on voltage stability and propose control strategies to mitigate voltage variations. The study will be conducted using Matlab, and the students will develop voltage control schemes that improve voltage regulation of power system to accommodate a high level of renewable energy. The students will be provided with appropriate literature and software tools as a start.

Project E2: Frequency Control of Renewable Energy Integrated Power System

Supervisor: Qianwen Xu, qianwenx@kth.se Electric Power and Energy Systems

The transition to renewable energy sources (RES) such as solar photovoltaics (PV) and wind power introduces significant challenges for frequency control in power systems. Traditionally, frequency stability has been maintained through the kinetic energy stored in the rotating masses of synchronous generators. However, renewable energy sources, particularly solar and wind power, do not inherently provide the inertia that conventional thermal and hydro generators offer, which complicates the control of frequency.

With the increasing share of renewable energy, the system's overall inertia is reduced, making it more vulnerable to frequency deviations during disturbances. Without adequate control, such deviations can lead to instability, triggering protection systems, load shedding, or even blackouts.

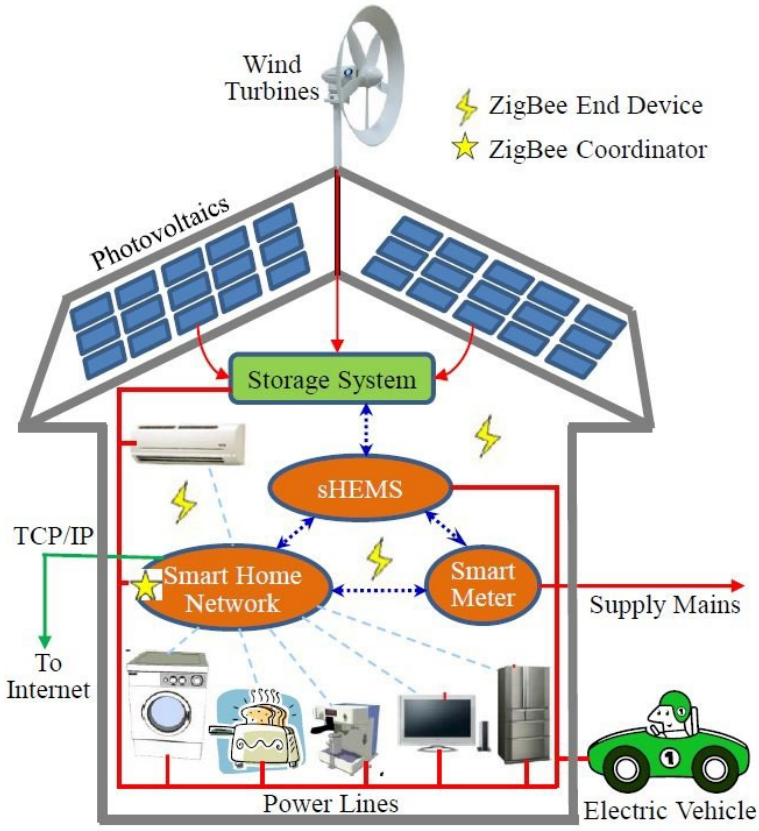
This project aims to investigate the challenges of frequency control in a renewable energy-dominated power system, and explore possible solutions, such as using battery energy storage systems (BESS), and developing advanced control strategies for inverter interfaced RESs. The students will be provided with appropriate literature and software tools as a start.

Project E3: Development of a future residential microgrid

Supervisors: Qianwen Xu, qianwenx@kth.se, Zhipeng Li zhipli@kth.se Electric Power and Energy Systems

Driven by environmental concern and sustainable requirement, development of residential microgrids attracts much attention around the world, as a forward step towards future carbon-neutral society. A residential microgrid is a small power system for a house/building, which consists of a solar photovoltaic (PV) source, a battery storage and residential loads, and can operate either in isolation or in connection to the main grid. In the daylight, the solar PV source can generate electricity to supply the loads, and the extra electricity can be stored in the battery to be used in the evening, or even sold back to the main grid. Thus a residential microgrid can reduce the energy cost and reduce CO2 emission. To make it works, each component of the microgrid should be properly designed and they should be controlled in a coordinated manner to provide stable and sustainable electricity

This project will develop a residential microgrid and its control scheme to achieve stable and sustainable electricity supply. The PV converter system will be designed to maximize its power generation in the daylight; the battery converter system will be designed to be charged when there is surplus electricity, and discharged when there is insufficient electricity. A coordinated control scheme will be developed for the whole system with high reliability and stability. The microgrid system will be developed in Matlab/Simulink as a demonstration of this project. The students will be provided with appropriate literature and some basic Matlab/Simulink models as a start. The students also have chance to deploy the solution in the microgrid hardware platform if time allows.



Project E4: Estimation and forecasting of the electricity prices

Supervisor: Mohammad Reza Hesamzadeh, mrhesa@kth.se,
Electric Power and Energy Systems

In this project we focus on estimating and forecasting electricity prices in the wholesale electricity markets. We mainly focus on the electricity spot markets which are day-ahead, intra-day, and real-time markets. Due to competitive forces in the today's electricity markets, electricity-price estimation and forecasting has become a fundamental tool which provides input to the decision-making mechanisms.

The electricity as a tradable commodity is notoriously volatile. This is partly because electricity is not economically storable and what is produced at a moment must be consumed at that moment. Also, electricity demand depends on hard-to-predict parameters such as weather or the intensity of everyday activities. These characteristics of electricity make the electricity prices to have a very complex dynamic. It often depends on several driving factors.

Besides, the current push from governments in many jurisdictions to increase the share of renewable CO2 free generating technology (mainly wind and solar generation) has added extra level of complexity to electricity price estimation and forecasting models. These renewable generation sources are intermittent, and they make the electricity spot prices more volatile than before.

At this background, the area of price estimation and forecasting has been quite active over the last few years. Various estimation-and-forecasting techniques are suggested in the academic literature with various degrees of success. Also, competition events such as Global Energy Forecasting Competition (GEFcom) are organized to attract the innovative forecasting techniques.

Broadly speaking, the estimation-and-forecasting techniques can be categorized as regression models and intelligent models. Linear and nonlinear regression are examples of regression models and the Neural network models are examples of intelligent models. The estimation and forecasting models can be static (without considering time) and dynamic (with time consideration). References [1] and [2] below provide very good information about different estimation-and-farecasting techniques.

At this background, this project has the following aims:

(1) A review of different recent techniques developed and suggested in the literature for electricity spot-price estimation and forecasting; In this review, the strengths and weaknesses of these reviewed techniques are clarified.
(2) To select three promising forecasting techniques and justify why these technique are suitable for forecasting; These three techniques can be selected between the regression and intelligent techniques.
(3) To apply the selected three forecasting techniques to estimate-and-forecast the spot prices in the Nordic electricity market. You can select two markets out of three following markets: day-ahead market, intra-day market, and real-time market. For your application, you may use the following software packages: Julia/Python, R, Eviews or Matlab.
(4) To interpret and explain the estimation-and-forecasting results that you have obtained for your selected spot markets and your forecasting technique.

References:
[1] https://en.wikipedia.org/wiki/Electricity_price_forecasting
[2] Bunn, Derek W. "Modelling prices in competitive electricity markets." (2004).

Context F: Hydro power modelling



A fundamental property of a power system is that the generation and consumption of electric power must always be in balance. It is therefore necessary to have enough flexible resources (i.e., generation or demand that can be adjusted to help maintain the balance of the power system) available. The amount of weather dependent, continuously varying generation (for example wind or solar power) is expected to continue increase in the future, which means that the need for flexible resources will also continue to increase.

An individual hydro power plant is very flexible and can rapidly increase or decrease the generation. However, the operation of multiple hydro power plants in the same river system will be most efficient if coordinated, as water released from an upstream reservoir will eventually reach the next reservoir in the river and if that reservoir is full then water will have to be spilled. The operation planning of a river system can be formulated as an optimisation problem, where the objective is to maximise the value of the total hydro generation in the river system, while taking into account the hydrological coupling between the hydro power plants as well as other operational limitations.

Hydro power has been one of the main sources of flexibility in the Nordic power system. However, to fulfil EU-wide goals for water environment, Sweden has a national plan for revision of the hydro power plant licences, which determine how the owners of the hydro power plants

are allowed to schedule generation. The revision of the licenses will need to balance conflicting environment goals and energy goals. For example, requirements on keeping water flows through the natural riverbeds will promote biodiversity and provide ecosystem services but will result in lower hydro power generation and may influence the flexibility of the concerned hydro power plants.

Large-scale energy system models are necessary to study possible solutions for a future carbon dioxide free energy system in Europe. In such models, it is vital to correctly model the flexibility of hydro power as hydro reservoirs are by far the largest energy storages in the system (for example, the Norwegian hydro power reservoirs can store 80 TWh). It is not possible to model individual hydro power plants in a large-scale energy system model. The solution is then to aggregate all hydro power in an area into a so-called equivalent mode. The idea is that this equivalent should be able to mimic the behavior of a detailed model while requiring significantly less computation power. The challenge is how to set up equivalent models to minimize the errors compared to a detailed model.

The aim of this context is to study how hydro power models can be improved and applied in practice. The operation planning of selected river systems will be studied using the open-source energy system modelling tool Spine [1].

Project F1: Environment conditions in River Ångermanälven

Supervisor: Mikael Amelin, amelin@kth.se, Electric Power and Energy Systems
The aim of this project is to set up a model of River Ångermanälven in the open-source software Spine and to run simulations to study the impact of different environmental conditions such as minimum flow through natural river beds and fish passageways.

Project F2: High-value production and high-value ratio

Supervisor: Mikael Amelin, amelin@kth.se, Electric Power and Energy Systems
The flexibility of a power plant is generally understood as the capability of the power plant to generate
electricity when it is needed by the system. However, there is no common definition of how to estimate
the flexibility of a power plant. This project will set up a small test system and investigate a newly suggested measure of hydro power flexibility [2] and compare it to earlier used flexibility measures such as flexibility factors, extreme load following capacity and equivalent energy storage.

Project F3: Equivalent models of open-loop pumped hydro

Supervisor: Mikael Amelin, amelin@kth.se, Electric Power and Energy Systems
An open-loop pumped hydro power consists of two reservoirs in a river system, where it is possible to pump water from the downstream reservoir to the upstream reservoir. It is challenging to set up an equivalent model of river system with regular reservoir hydro power as well as one or more pumped hydro power plants. This project will set up a small test system and compare the results of a detailed model to four different designs of the equivalent model.

[1] https://github.com/Spine-tools

[2] J. Bladh & J. Funkquist, ”Högvärdesproduktion och högvärdeskvot”, Energiforsk rapport 2026-1170, Energiforsk, Stockholm 2026.

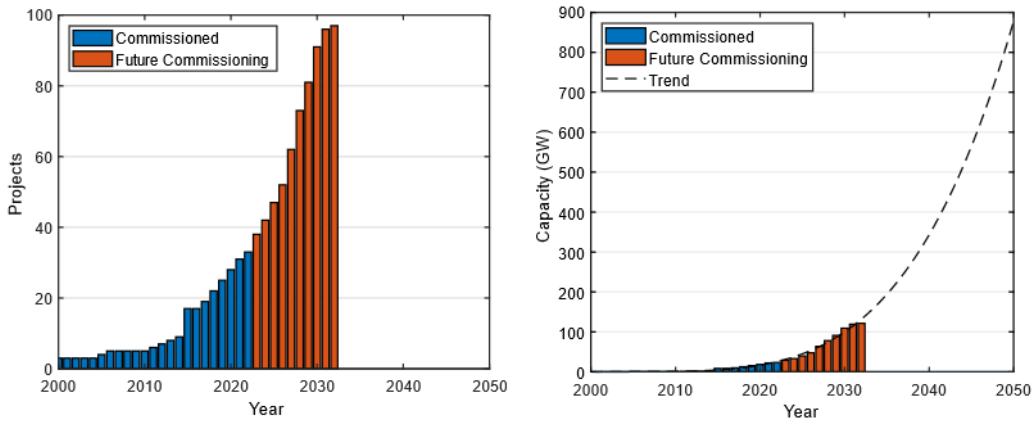
Context G:
HVDC for Supergrids & Offshore Wind

In Europe, we target the massive integration of offshore wind power into our power system. End of 2022, ca. 30 GW capacity was installed. The EU goal for 2050 is 300 GW [1]. Offshore wind is important to reach our CO2 targets and energy independency. Most offshore windpower will be connected using high-voltage direct-current (HVDC) technology. Today, many point-to-point HVDC connections exist, but no HVDC grids. However, the rapid development of HVDC supergrid technology is ongoing, and the Swedish industry is a world market leader in HVDC technology.

In this context, we will learn about different aspects of a North Sea Supergrid, offshore wind, and the technology “HVDC”.



Picture of an offshore HVDC station (source: PROMOTioN/TenneT)



(a) Cumulative projects by year (b) Cumulative installed capacity

Ongoing and expected HVDC boom [2]

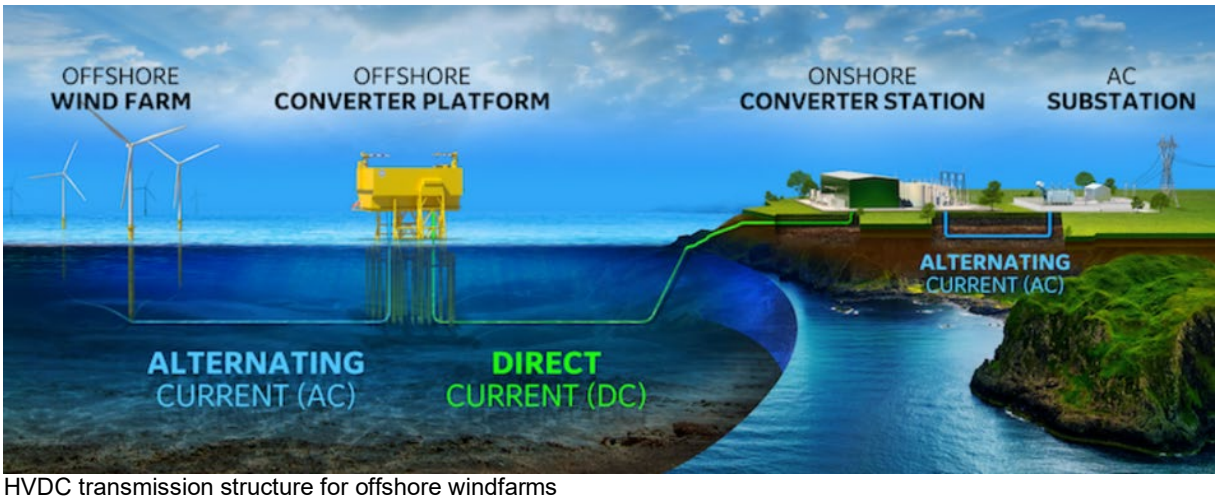
[1] Ostend Declaration of Energy Ministers On The North Seas as Europe’s Green Power Plant Apr. 24, 2023. Available online: https://www.government.nl/documents/diplomatic-statements/2023/04/24/ostend-declaration-on-the-north-sea-aseuropes-green-power-plant

[2] J. Cabañas Ramos, M. Moritz, N. Klötzl, C. Nieuwenhout, W. Leon Garcia, I. Jahn, and A. Monti, “Getting ready for multi-vendor and multi-terminal HVDC technology”, *MDPI energies*, vol. 17, no.10, pp. 2388-2416, May 2024. https://www.mdpi.com/1996-1073/17/10/2388

Project G1: Connecting offshore windfarms via HVDC compared with HVAC transmission

Supervisors: Zhenlin Zhao (zhenzhao@kth.se), Electric Power and Energy Systems

Offshore windfarms generate alternating current (AC) at wind turbines, but how to transmit electricity from the AC collector arrays to the onshore power grid is one of the most consequential engineering choices in offshore wind projects. The choice between high-voltage alternating current (HVAC) and high-voltage direct current (HVDC) transmission affects offshore substation design, cable specifications, transmission losses, and finally the efficient and reliable operation of offshore windfarms.



HVDC transmission structure for offshore windfarms

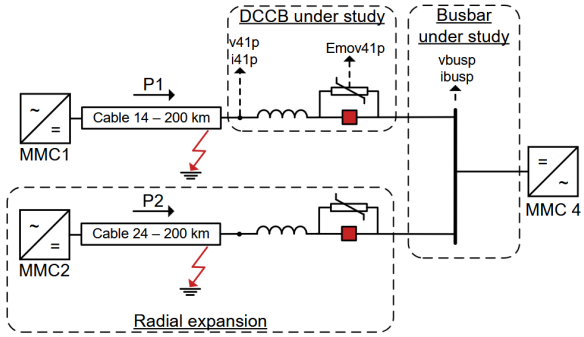
In this project, the students will learn:

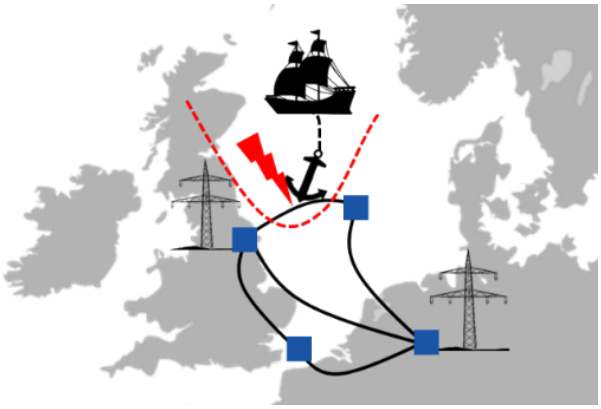
- How HVDC and HVAC transmission systems for offshore windfarms work.
- How to simulate the HVDC and HVAC systems for offshore windfarms in PSCAD.
- How to monitor and analyze the operation of offshore windfarms.
- How to quantify system performance, such as transmission losses, transmission efficiency, voltage stability, and other relevant performance indicators.
- How to compare HVDC and HVAC transmission systems and identify the benefits of HVDC transmission for offshore windfarms.

Project G2: Protection of expanding HVDC systems against blackouts

Supervisor: Ilka Jahn, ilka@kth.se, Electric Power and Energy Systems

Any electrical system can fail. In the power system, failures can originate from human error or natural disaster, for example, a ship anchor destroying a subsea cable. HVDC grids need to be designed to handle failures and contain possible damage to a certain grid part. A cascading blackout affecting several surrounding countries is not an option. In this project, we will investigate ways to protect HVDC systems, in particular, the question of whether already installed protection equipment (DC circuit breaker) is still powerful enough to protect the grid, even if the system is being expanded.





Case study in the KEX project (left), example cable fault in an HVDC supergrid (right)

In this project, the students will learn:

- About blackouts and protection in HVDC systems (e.g., anker on cable)
- How to simulate HVDC technology in PSCAD (based on existing simulation)
- How to simulate and investigate short circuits in computer simulation
- Expand a point-to-point HVDC system into a three-terminal system
- Investigate stresses on DCCBs before and after expansion

References

I. Jahn, N. Johannesson and S. Norrga, "Survey of methods for selective DC fault detection in MTDC grids," *13th IET International Conference on AC and DC Power Transmission (ACDC 2017)*, Manchester, UK, 2017

Project G3: Influence of HVDC system parameters on the dynamic performance of HVDC converters

Supervisor: Gerald Cham Kpu, geraldck@kth.se, Electric Power and Energy Systems

HVDC systems enable large-scale power transmission over long distances and interconnection of power systems. The dynamic behaviour of an HVDC converter depends not only on its controls but also on the electrical characteristics and configuration of the surrounding DC system. Components such as transmission lines and cables, DC reactors, capacitors, and grid configuration determine the impedance seen by the converter and influence its voltage and current dynamics. Changes in cable length or type, reactor size, capacitance, or DC-grid configuration can therefore significantly affect converter performance and stability through interactions between the converter controls and the DC system. In this project, we will investigate how these HVDC system parameters affect the dynamic performance of an HVDC converter. The students will simulate an HVDC system in PSCAD and vary parameters such as cable length, DC reactor, DC capacitance, etc.

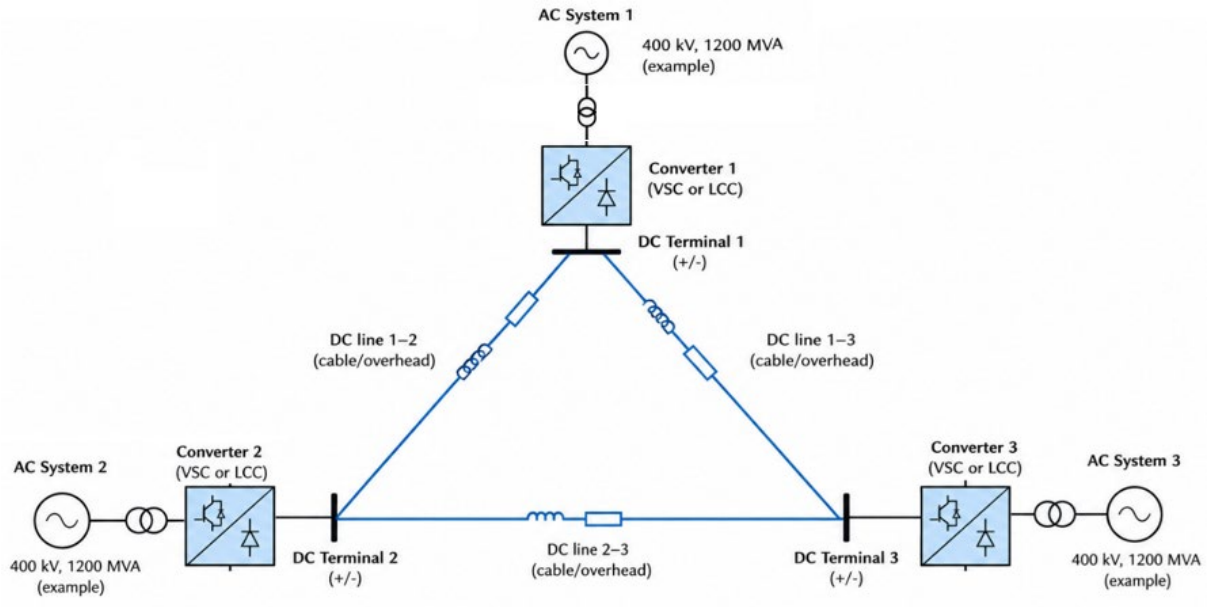


Figure 1: Example HVDC system

In this project, the students will learn:

- About HVDC transmission systems and HVDC converter performance dynamics
- How to perform HVDC simulations in PSCAD (based on an existing simulation)
- How to simulate and investigate how HVDC cables, DC reactors, capacitors, etc. affect the dynamic performance of an HVDC converter
- How to use data-analysis methods to investigate system parameter interactions and visualise the results.

References

P. Torwelle, K. Shinoda, P. Verrax, et al., “Functional requirements for HVDC grid systems and subsystems,” Interopera, 2025.

Context H: Design and testing of novel microwave/antenna technologies

Today, microwave technology is employed in many of our technological devices, and they fulfil an essential function in communication systems, intelligent cities, surveillance, medical diagnosis and space observation.

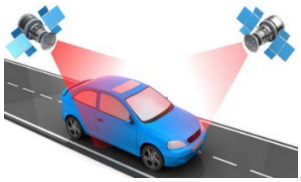
Innovative microwave designs are required daily in the products of technology-driven companies. These companies require efficient and multi-functional antennas and microwave devices that can enable:

- High data rate communications for present 5G and future 6G networks.
- Efficient satellite communications with the newly deployed low-Earth-orbit satellites.
- High resolution radars to detect people, vehicles and objects in smart cities.
- Non-invasive imaging of patients for early detection of health issues.
- Highly precise airport scanners that maximize the location of concealed objects.
- Precise detection of stars and planets in outer space.

Within the projects of this Context, you will be able to acquire the fundamental knowledge for designing advanced microwave devices and antennas. You will learn how to use commercial software of simulation, which is commonly employed in industry. Finally, you may manufacture and measure a proof-of-concept. After the project, you will be able to reproduce the usual steps followed in a microwave or antenna design process.

Examples of challenges which require innovative microwave/antenna technologies:



Picture 1. Artistic rendition of future communications.
Picture 2. Autonomous car inter-connected with wireless systems.
Picture 3. Patient inside a high-resolution medical scanner.
Picture 4. Car communicating with low-Earth-orbit satellites.

Project H1: Study on modeling of electromagnetic radiation of planar phased array and radome with ray tracing

Supervisors: Oscar Quevedo-Teruel, oscarqt@kth.se, Riccardo Musso rmusso@kth.se, Núria Flores Espinosa nuriafe@kth.se, *Department of Communications Systems*

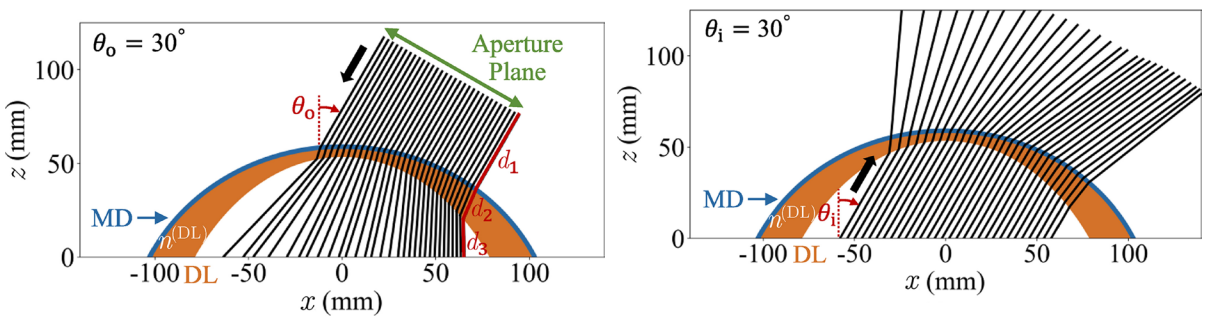
Phased arrays are widely used in modern radar, sensing, and communication systems due to their capability to electronically control the direction and shape of radiated beams. When a phased array is integrated with a radome, the electromagnetic interaction between the array, the radome, and the surrounding environment can significantly affect its radiation and scattering properties. Therefore, accurate and efficient methods for analyzing electromagnetic scattering are important for the design and characterization of these systems.

Full-wave electromagnetic simulations can provide highly accurate results, but their computational cost can become significant when analyzing electrically large structures or evaluating scattering over many observation angles. Ray-tracing techniques provide an efficient alternative by modeling the propagation of electromagnetic rays that can be particularly useful for estimating the radiation pattern (RP) or the radar cross section of complex structures.

This project will be aimed at developing a two-dimensional ray-tracing (2D-RT) model for the calculation of the RP of a planar phased array integrated with a radome. The model will consider the interaction of incident rays with the radome and the phased-array structure and will be used to calculate the overall pattern considering the total contribution from each element of the array. The results will be compared with full wave electromagnetic simulations to assess the accuracy and limitations of the ray-tracing approach.

As a student, you will learn:

- The fundamentals of ray tracing, electromagnetic scattering and calculation of the radiation pattern (RP).
- How to implement a simplified 2D ray-tracing algorithm, comprising of both “direct” and “inverse” ray tracing problems.
- Calculation of reflection and transmission through dielectric structures.
- Calculation and visualization of RP as a function of the observation angle.
- Comparison between ray-tracing and full-wave electromagnetic simulations.



Some examples of direct and inverse 2D-RT of radome. Extracted from H. Wang, M. Chen, F. Mesa, and O. Quevedo-Teruel, “ Hybrid Dielectric–Metadome Structures Modelled via 2-D Ray Tracing and Physical Optics,” *IET Microwaves, Antennas & Propagation* (2026).

Project H2: Fully metallic geodesic lens with circular polarization conversion

Supervisors: Oscar Quevedo-Teruel, oscarqt@kth.se, Dayan Pérez-Quintana dayanpq@kth.se, *Department of Communication Systems*

The growing demand for higher data rates, network capacity, and low latency in future wireless systems is driving the adoption of millimeter-wave and sub-THz frequencies for 5G and 6G applications. Although these frequency bands provide large bandwidths, they also suffer from increased propagation and material losses, requiring highly efficient high-gain antenna solutions.

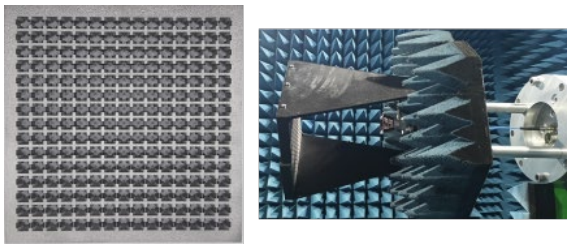
Quasi-optical beamformers (QOBFs) have recently emerged as a promising technology for advanced antenna solutions in both mobile and satellite communications. Most existing QOBFs are implemented using parallel-plate waveguide (PPW) technology and operate in monomodal regime, enabling low-loss beamforming and compact architectures. However, due to their inherently monomodal operation, these systems typically support only a single polarization, which limits their applicability in many practical communication scenarios.

Modern communication systems often rely on dual-linear or circular polarizations to improve capacity and enable satellite communications. However, polarization diversity in PPW-based beamforming architectures has received limited attention in the literature. Enabling polarization diversity in low-loss beamforming networks is particularly relevant for future satellite systems and emerging non-terrestrial 6G networks.

This project proposes a low-profile fully metallic geodesic lens for linear-to-circular polarization transformation in the Ka-band. Based on a PPW polarizer with broadband matching and efficient polarization conversion, the design seeks to improve the performance of next-generation terrestrial and satellite communication systems. The work will involve electromagnetic simulations, device optimization, fabrication, and experimental characterization.



(a) Fully metallic geodesic Luneburg lens and integrated PPW polarizer implementation.



(b) Fully metallic polarizer prototype and experimental characterization setup.

As a student, you will learn:

- Fundamental principles of antenna polarization and polarization conversion techniques.
- How to simulate periodic structures and quasi-optical beamforming devices using commercial electromagnetic software.
- How to design and implement fully metallic geodesic lenses and PPW-based polarizers for broadband operation.
- How polarization performance is experimentally characterized and what fabrication technologies are available for fully metallic devices.

Project H3: Design of a matched, flat graded index lens at millimetre-wave frequencies

Supervisors: Oscar Quevedo-Teruel, oscarqt@kth.se, Moises Tercero moisestn@kth.se
Department of Communication Systems

To keep up with the rising needs of higher data rates, wireless communication systems are moving to higher frequencies in the millimetre wave spectrum to obtain broader bandwidths. However, as frequency increases, path loss becomes more significant, so every ounce of power must be leveraged. For this reason, high-gain antennas are desirable to compensate for the reduced link margin.

A way to implement highly directive antennas is to build antennas with large apertures such as lens antennas. However, conventional (curved) dielectric lenses can be thick and difficult to manufacture in contrast to flat lenses. To make flat lenses, the refractive index distribution across the aperture must be varied. This is achieved, for example, by drilling holes of varying size in a dielectric disc. However, reproducing a given refractive index profile is not enough to make the most out of the available power; some power is lost due to reflections at the air-lens interface. While reflections in conventional lenses can be mitigated through quarter-lambda matching sections, the impedance matching of graded refractive index (GRIN) lenses is more complicated due to their inhomogeneity.

This project aims to design, manufacture and experimentally characterize a matched flat GRIN lens. In particular, the project focuses on the development of a simple impedance matching strategy that results in flat and thin anti-reflection (AR) layers.

As a student, you will learn:

- Fundamentals of aperture antenna design
- How to model heterogeneous arrangements of air and dielectrics
- How to model the scattering of multilayer structures under oblique illumination
- How to leverage the model to produce impedance-matched lenses
- How to simulate a lens antenna and carry out its measurement

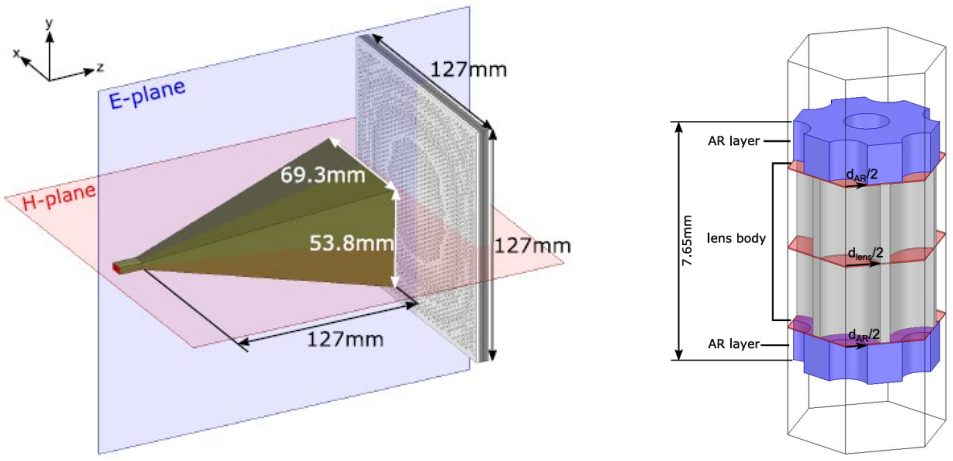


Figure: (Left) Simulation model of a horn antenna with a dielectric GRIN lens. (Right) Unit cell of the lens figuring a multilayer dielectric stack. Pictures obtained from [2].

[2] Y. He and G. V. Eleftheriades, "Matched, Low-Loss, and Wideband Graded-Index Flat Lenses for Millimeter-Wave Applications", in *IEEE Transactions on Antennas and Propagation*, vol. 66, no. 3, pp. 1114-1123, March 2018.

Project H4: Design of a waveguide feeding network for satellite communications using Electromagnetic Bandgap (EBG) periodic structures

Supervisors: Oscar Quevedo-Teruel, oscarqt@kth.se, Hairu Wang hairu@kth.se
Department of Communication Systems

With the development of modern society and the increasing demand for high-performance satellite communications, robust and efficient antenna systems are essential for satellite payloads. These systems typically consist of both the antenna and its feeding network.

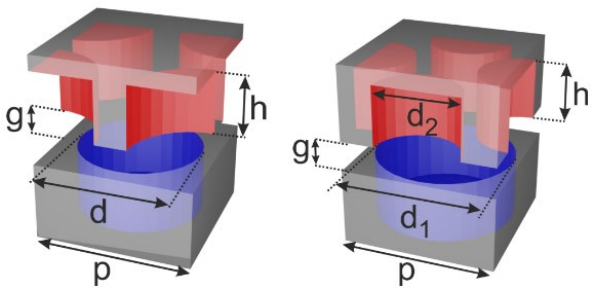
Feeding networks are guided-wave structures used to distribute electromagnetic power in antenna systems. They are often assembled from several metallic plates, where small misalignments or air gaps may cause electromagnetic leakage and increased transmission loss, especially at high frequencies.

In recent years, periodic structures have been widely investigated for the design of various microwave components, including lenses, leaky-wave antennas, and low-leakage waveguide interfaces and flanges. Different periodic structures exhibit different electromagnetic properties, and selecting an appropriate structure is therefore crucial for a given application.

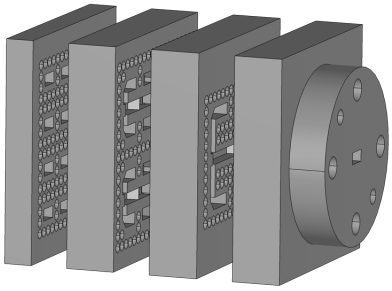
This project investigates holey electromagnetic bandgap (EBG) structures for suppressing leakage in waveguide feeding networks. In particular, a glide-symmetric holey EBG and a novel three-hole EBG will be compared in terms of leakage suppression, bandwidth, and compactness. The work will include electromagnetic simulation and design, and, if feasible, fabrication and experimental characterization of the feeding network.

As a student, you will learn:

- The electromagnetic principles and operation of periodic structures.
- How to simulate and analyze periodic structures using commercial electromagnetic software.
- How to design and simulate waveguide feeding networks using commercial electromagnetic software.
- How to experimentally characterize and measure the performance of a waveguide feeding network.



Unit cells of Holey EBG periodic structures



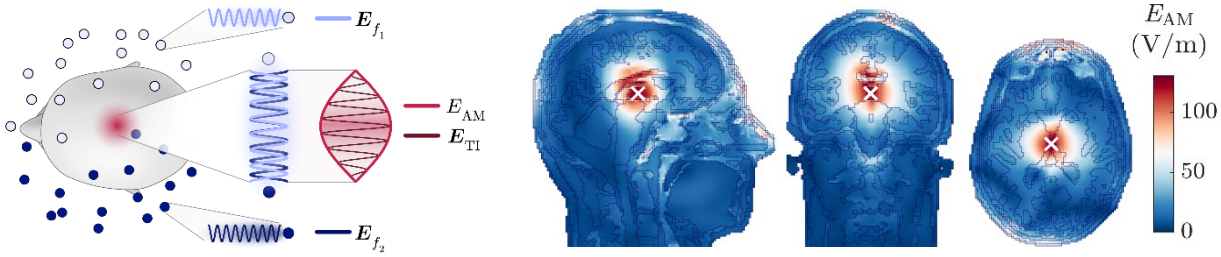
A waveguide feeding network

Project H5: Microwave focusing in a human brain model for deep brain stimulation

Supervisors: Mariana Dalarsson, mardal@kth.se, Mika Söderström, mikaso@kth.se
Department of Electromagnetics and Plasma Physics

Stimulation of biological tissues using electromagnetic (EM) fields is an active area of research within biomedical engineering. One important application is deep brain stimulation (DBS), used to treat neurological disorders such as Parkinson's disease, epilepsy, and other brain disorders. Current DBS treatments rely on surgically implanted electrodes that directly stimulate neural tissue. While effective, these procedures are invasive and involve risks associated with surgery and implanted devices. As a result, there is significant interest in developing non-invasive techniques capable of stimulating deep brain regions without affecting surrounding tissue.

Recent research within our group investigates microwave-based DBS using arrays of external antennas. By carefully controlling the signals applied to the antennas, EM fields can interfere constructively in selected deep brain regions while remaining weaker in surrounding areas. This emerging approach has the potential to provide highly localized stimulation without the need for implanted electrodes.



Figures: Microwave focusing in a human head model. Left: Antenna elements (dots) surrounding a head model. Right: Example of a focused EM field distribution obtained in previous work within this research group [1].

The project is part of an ongoing research effort aimed at understanding and improving microwave-based techniques for non-invasive DBS. The objective of this project is to gain an understanding of how EM fields can be focused inside realistic human head models, as well as to further develop existing simulation methods for this application. The findings and insights gained from this project will support future research in microwave-based DBS.

In this project, the students will learn:

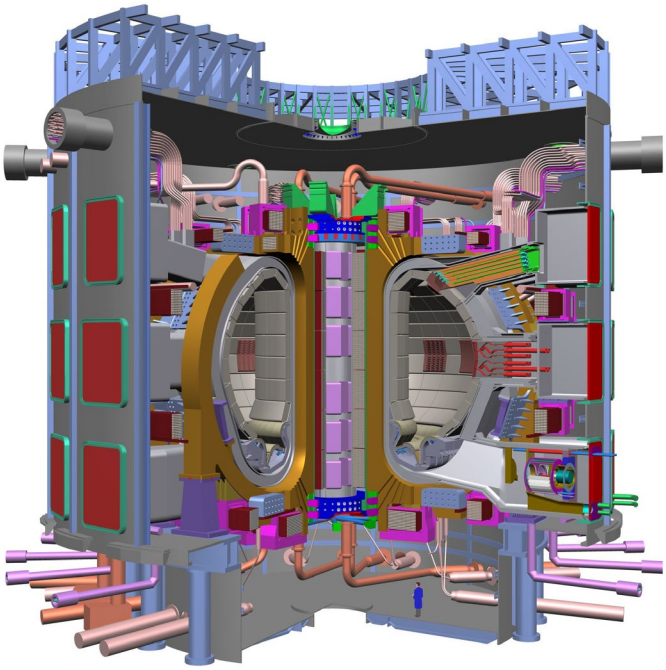
- Fundamentals of deep brain stimulation and neuro-modulation.
- Electromagnetic wave propagation in biological tissues, including principles of temporal interference and electromagnetic focusing.
- How to use commercial simulation software for biomedical EM applications, including how human-body EM simulations are constructed, validated, and interpreted.
- How to evaluate simulation results using performance and safety metrics.

[1] M. Söderström, et.al, "Microwave Focusing with Temporal Interference for Non-Invasive Deep Brain Stimulation," in *IEEE Transactions on Antennas and Propagation*, doi: 10.1109/TAP.2026.3722231

Context I:
Fusion, solens energikälla på jorden

Introduktion

Fusionsforskningen arbetar för att kunna konstruera ett kraftverk som genererar energi från kärnreaktioner mellan olika väteisotoper. Dessa fusionsreaktioner avger ungefär en miljon gånger mer energi än kemiska reaktioner och är den process som värmer solen. Om fusionskraften kan bemästras på jorden har vi en i princip i outtömlig energikälla utan växthuseffekter och med relativt lite radioaktiva restprodukter. Liknande reaktioner sker dagligen i fusionsexperiment världen över, men man har aldrig lyckats producera mer än 65% av den inmatade effekten. För att producera nettoeffekt krävs större experiment och just nu byggs en experimentanläggning, ITER, i södra Frankrike som väntas producera tio gånger högre effekt än vad man stoppar in. ITER är det andra mest påkostade vetenskapliga projektet i mänsklighetens historia, efter den internationella rymdstationen (ISS). Om fusion fungerar, som många forskare tror, kan det ha stor betydelse för vår framtida energiförsörjning.



Den 73m höga fusionsreaktorn ITER som nu byggs i södra Frankrike och som ska stå klar 2020. ITER väntas kunna producera 10 gånger mer energi än den förbrukar.

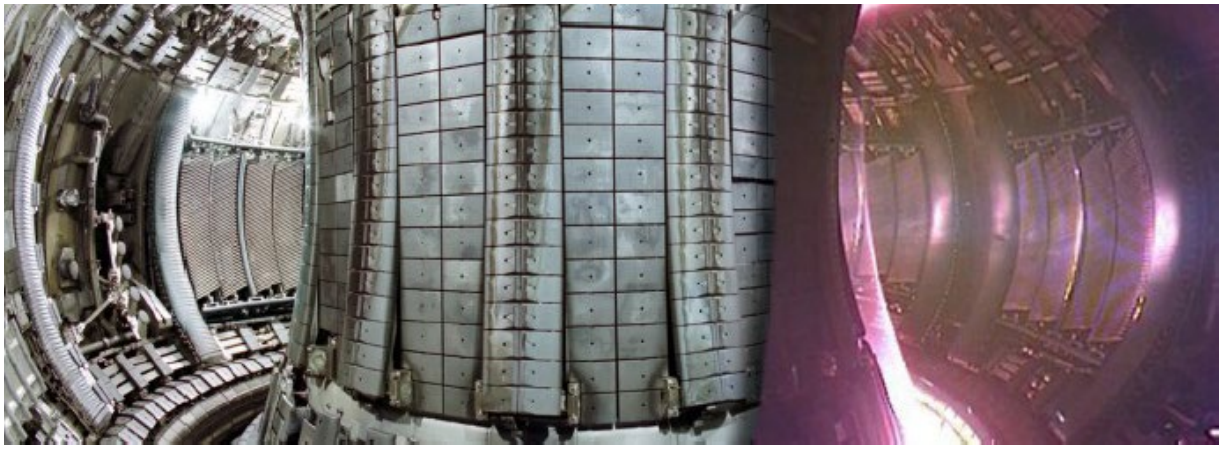
Varför behövs så stora experiment? För att fusionsreaktionen ska komma i gång måste man uppnå en temperatur på över 200 miljoner grader, samtidigt som man behöver en tillräckligt hög täthet (~10^{20} m^{-3}), vilket är svårt att åstadkomma i mindre maskiner. Så hur kan man bygga en reaktor som innesluter en 200 miljoner grader varm gas (eller plasma som gasen kallas vid dessa temperaturer)? Det finns inga material som klarar att värmas till över 3 000 grader utan att smälta, så i en fusionsreaktor måste det varma plasmat hållas borta från väggarna. Detta sker med hjälp av magnetfält. Men även med starka magnetfält "läcker" värmen ut, och väggarna i en fusionsreaktor utsätts för stora påfrestningar. Dessutom måste plasmat ständigt värmas upp för att kompensera för värmeförluster till väggen. Denna uppvärmning kommer dels från fusionsreaktionerna, dels från injektion av radiovågor och av högenergetiska partiklar, samt resistiv uppvärmning.

Projekten i den här kontexten ingår både experimentella och teoretiska projekt. Här får man möta forskning vid frontlinjen och man får en inblick i möjligheterna och utmaningarna kring fusionskraften. Dessutom ska vi besöka fusionsexperimentet Extrap-T2R på KTH, samt diskutera etiska och politiska frågor kring vår framtida elförsörjning.

Projekt I1: Calibration and interpretation of ion beam measurements performed on walls in fusion machines.

Supervisor: Per Petersson, ppeter@kth.se

Under the conditions that apply on earth (gravity, air pressure etc.), hot fusion plasmas must be performed in a vacuum vessel and confined by strong external forces. Despite the magnetic confinement the wall surrounding the plasma both gets eroded and have material deposited on top of it. These deposits contain most elements present in the machine and is sometimes referred to as “tokamakium”. It can have very different content and properties depending on where and how it is formed. In future reactor devices, such as the International Thermonuclear Experimental Reactor (ITER), the interaction of the plasma with surrounding materials in the vacuum vessel constitutes one of the main remaining engineering problems.



JET tokamak in Culham, England - with and without plasma that interacts with the wall. Visually you can see deposits in darker areas on the inner wall guard limiters that stick out from the central column, but deposits of different thicknesses are present on most surfaces inside the machine if you look hard enough.

An important family of tools that we frequently use for looking at tokamakium, as well as other very different samples, is ion beam analysis, that uses accelerators that generates ion beams used to look at materials. One specific technique is Time of Flight Elastic Recoil Detection Analysis (ToF-ERDA). In the analysis of the data from ToF ERDA several steps of manual interpretation are required for good results, especially in some tricky situations such as separating isotopes close in mass such as ^{11}B and ^{12}C, that both can be present in tokamakium.

The main goal in this project will be to use both simulated data and real data to compare traditional and automated methods for analysis of ToF ERDA data. To get a good set of data we start by manufacturing both simple and more complicated reference samples by magnetron sputtering and then perform ToF-ERDA measurements. The data produced can be used both for calibration and as test of the result. One technique that can be tested for interpreting the data is some form of machine learning, but other options are possible such as extending the method for energy calibration.

As part of the project, you will be introduced to accelerator-based material analysis techniques, sputtering techniques to build materials and data analysis with and without machine leaning. The work will include visits to the Tandem Accelerator Laboratory of Uppsala University that has the necessary accelerator to perform the measurements.

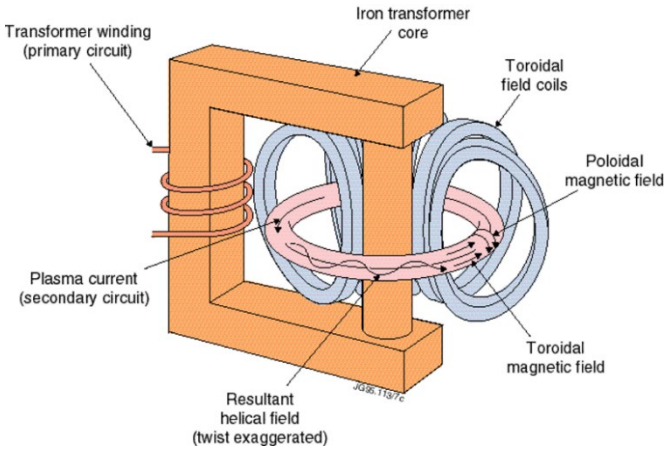
The project is part of larger program of investigation of material and can be adjusted in cooperation with the student in several directions such as to better understand the background noise in the measurements.

Project I2: EXTRAP T2R Iron Core Stray Magnetic Field

Supervisor: Per Brunsell, brunsell@kth.se, EMP

Background

A common device in fusion energy research is the torus-shaped plasma discharge magnetically confined inside a set of toroidal magnetic field coils. A toroidal electrical current is formed in the plasma to produce a poloidal magnetic field component, which gives the desired helical twist of the confining magnetic field. This magnetic confinement scheme is used in the EXTRAP T2R device at KTH.



The toroidal plasma current is produced by magnetic induction using an iron core transformer and a primary circuit connected to a capacitor bank power supply. To extend the maximum pulse length of the plasma, the iron core is pre-magnetized to an initial negative bias magnetic flux by a core bias primary winding.

An undesired side effect of the pre-magnetization of the iron core is the stray magnetic field produced in the plasma volume. This steady state stray magnetic field needs to be compensated by a set of saddle-coil type correction magnetic field coils, distributed at the plasma torus boundary.

Estimates of the local steady state field at the plasma torus are required as complement to the magnetic flux loop sensor measurements that are used for the feedback-controlled correction coils. The reason is that the flux loops sensors only measure the time-dependent component of the field.

Project goals

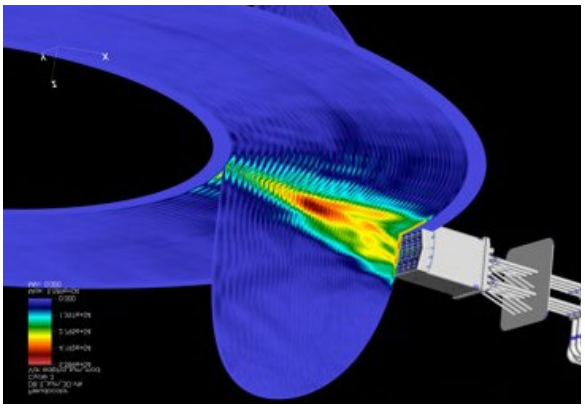
- The project goal is to develop estimates of the local steady state magnetic field produced by the iron core pre-magnetization using 3D finite element magnetic field calculations calibrated against magnetic field measurements on the EXTRAP T2R device.

- The calculation of the stray magnetic field is performed with the COMSOL modelling software, using a 3D model of EXTRAP T2R that incorporates the geometry of the iron core and the core bias windings.

- Measurements of the steady state magnetic field are done with a small set of Hall sensors that are available at EXTRAP T2R, sampling the stray field at a few locations. The COMSOL model will then be used to interpolate the field to all other positions at the plasma torus boundary.

Projekt I3: Modellering av radiovågsuppvärmning av fusionsplasma

Supervisor: Thomas Jonsson, johnso@kth.se, EMF; Lukas Bähner, bahner@kth.se, EMF



***Figur**: Elektrisk fältstyrka från en numerisk simulering av radiovågsuppvärmning i ett fusionsplasma.*

Bakgrund

Fusionsreaktioner kräver mycket höga temperaturer. För att producera stora mängder energi i ett fusionskraftverk behöver plasmat upphettas till temperaturer på omkring 200 miljoner grader. En av de mest använda metoderna för att värma plasmat till dessa höga temperaturer är att använda radiovågor.

Trots att radiovågsuppvärmning har använts inom fusionsforskningen i över 50 år upptäcktes för ungefär tio år sedan en ny metod, så kallad 3-jonuppvärmning (3-ion heating) [Kazakov et al., *Nucl. Fusion* 55 (2015) 032001]. I detta projekt ska vi studera denna metod med hjälp av numeriska simuleringar med kodpaketet FEMIC-Foppler. Resultaten från simuleringarna ska därefter jämföras med tidigare publicerade experimentella resultat från fusionsanläggningen EAST [Jin et al., *Phys. Plasmas* 27 (2020) 082506].

Projektets innehåll

I projektet kommer ni att lära er grunderna i hur en fusionsreaktor fungerar, hur plasma kan inneslutas med hjälp av magnetfält och hur plasma kan värmas med radiovågor. Ni kommer att genomföra egna numeriska simuleringar med FEMIC-Foppler, där FEMIC löser en vågekvation medan Foppler beräknar hur snabba joner kan accelereras av radiovågorna.

Ni kommer även att studera en vetenskaplig artikel som beskriver experiment och simuleringar från EAST, där en annan simuleringskod har använts. Resultaten från denna studie kommer sedan att jämföras med era egna simuleringar. Arbetet kommer huvudsakligen att genomföras i MATLAB, men även COMSOL Multiphysics kommer att användas.

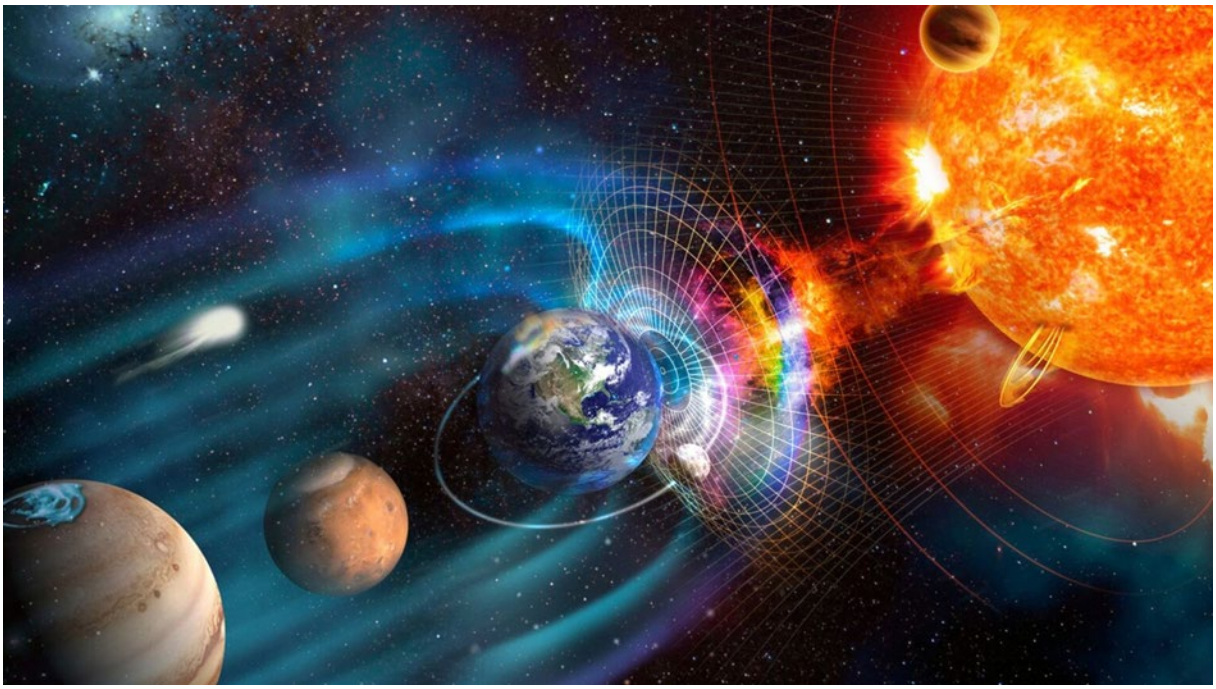
Projektets mål

1. Läsa in sig på grundläggande litteratur om fusionsplasmafysik och 3-jonuppvärmning.
2. Lära sig den grundläggande fysiken bakom FEMIC och Foppler samt att köra koderna och analysera resultaten.
3. Studera rapporten av Jin et al, som senare ska användas som grund för jämförelsen med de egna simuleringarna.
4. Utforma simuleringar som kan användas för att undersöka och demonstrera 3-jonuppvärmning.
5. Genomföra simuleringar som kan jämföras med experimentella resultat från EAST.
6. Jämföra simuleringsresultaten med resultaten från EAST och:
    a. studera skillnaderna i detalj,
    b. identifiera möjliga orsaker till skillnaderna, exempelvis skillnader mellan de matematiska modeller som har använts.
7. Skriva en vetenskaplig rapport och presentera projektets resultat muntligt.

Context J:
Solar wind and planetary environments

Context Responsible: Lorenz Roth (lorenzr@kth.se)
Space Physics, Department of Electromagnetics and Plasma Physics (EMP)

Space Physics encompasses the physics of the open space in our solar system, mainly the environments of the Earth and the Moon, other planets and moons, and the Sun. The neutral gas and plasma (charged particle) environments of the Sun, the planets (including Earth's magnetosphere) and smaller bodies like moons and asteroids are studied with the help of space probes that are in high-altitude orbit around the Earth or that visit other planets. Observations are also made by space-based telescopes, such as the James Webb Space Telescope (JWST). The space physics research group at KTH is involved in various projects that utilize direct in-situ measurement by space probes from both NASA and ESA missions as well as the observatories mentioned. In this context, students have the possibility to participate in real research projects within observational space physics.



Project 1 uses data from the NASA Juno spacecraft of investigate turbulence in the environment of our largest planet Jupiter.

Projekt 2 uses brand new data from the NASA/ESA James Webb Space Telescope to study variability in the outgassing from a small moon of planet Saturn.

Projekt 3 studies how the shape of turbulent structures in space plasmas affects energy dissipation applying Hall-MHD simulations and Schur decomposition.

Project 4 uses GNSS data from the Swedish SWEPOS network to study pulsating aurora, analyzing Total Electron Content variations to characterize their spatial and temporal properties and compare them with optical observations.

Project 5 studies how different lunar orbits and starting points affect the accuracy of orbit predictions by comparing simulations with real spacecraft data.

Project 6 uses data from the NASA Pioneer Venus Orbiter to statistically characterize Venus's upper atmosphere and ionosphere.

Project J1: Turbulent cascades in Jupiter's magnetospheric plasma disc

Supervisor: Anatol Große-Schware, anatolgs@kth.se, Electromagnetics and Plasma Physics

Jupiter's magnetosphere is filled with plasma that origins mainly from the volcanic moon Io. The plasma accumulates in a donut-shaped region that is corotating with the gas giant, the Io plasma torus. This torus is dipped against the rotational equator because of the axial tilt of Jupiter's magnetic dipole axis with respect to the rotation axis. At larger distances from Jupiter, the Io plasma torus diffuses into a magnetospheric plasma sheet. Within the corotational plasma flow in both, the torus and the sheet, turbulent phenomena occur. This project focuses on the turbulent cascade, i.e., the energy transport from large vortices to small vortices based on magnetic field fluctuations. These turbulent cascades are characterized by spectral indices that describe the power drop from large scales to smaller.

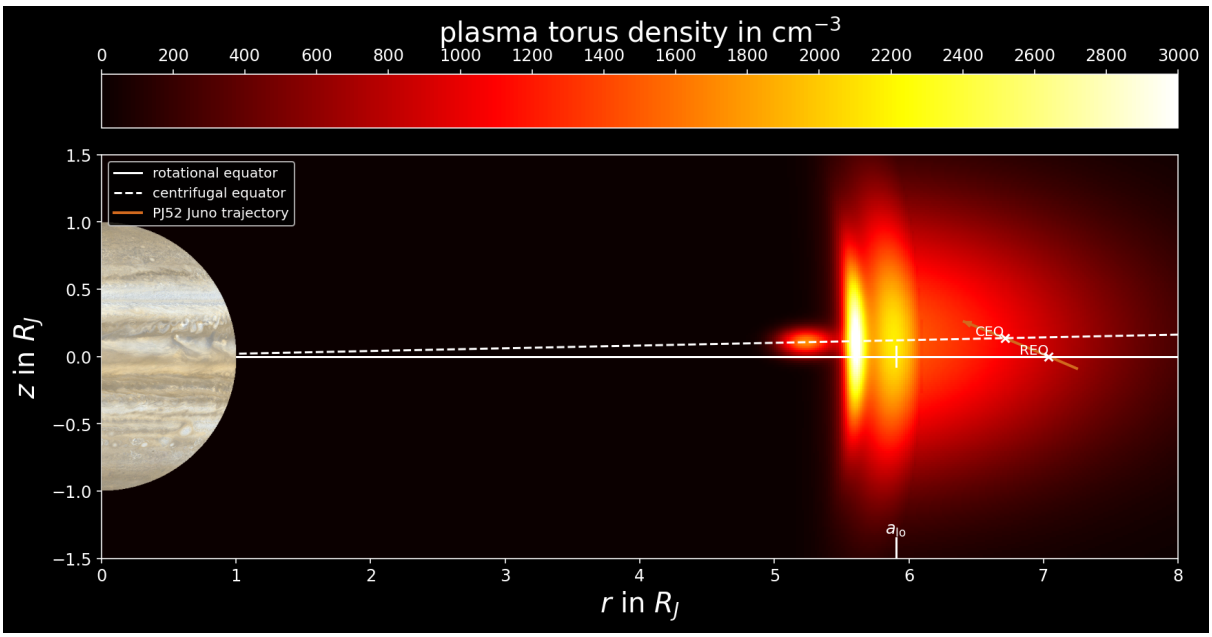


Figure 2: Trajectory of Juno's plasma torus crossing during the PJ52 orbit

In this project, students will use magnetic field data from the Juno spacecraft and investigate the variability of the turbulent cascade and their spectral indices with respect to the radial distance from Jupiter. We are interested in all Juno plasma sheet crossings that happened between the orbits of the moons Io (at 5.9 Jupiter radii) and Callisto (at 26.3 Jupiter radii).

Steps included:

1. Find the Juno magnetometer datasets that contain the respective plasma sheet crossings.
2. Perform a spectral analysis of the magnetic field data.
3. Determine the spectral indices of the turbulent cascade for both, MHD scales and kinetic scales.
4. Obtain a radial profile of the spectral indices.

By carrying out this project, students will contribute to ongoing space science research and gain experience in the analysis of spacecraft data, statistical methods and turbulent physics.

Project J2: New James Webb Space Telescope observations of water plumes from an icy moon

Supervisor: Lorenz Roth, lorenzr@kth.se, Electromagnetics and Plasma Physics

Saturn's icy moon Enceladus is a primary target for astrobiology because of its active cryovolcanic plumes, which eject material from a subsurface ocean into space. This plume contributes to a neutral gas torus along Enceladus' orbit and provides direct insight into the chemistry of the subsurface ocean and its potential habitability. Recent observations with the James Webb Space Telescope (JWST), using the Near Infrared Spectrograph (NIRSpec), have enabled remote detection and spatial mapping of water-vapor emission from near-Earth space with unprecedented sensitivity, as shown in the figure below. And in November and December 2026, JWST will observe Enceladus and its plume again!

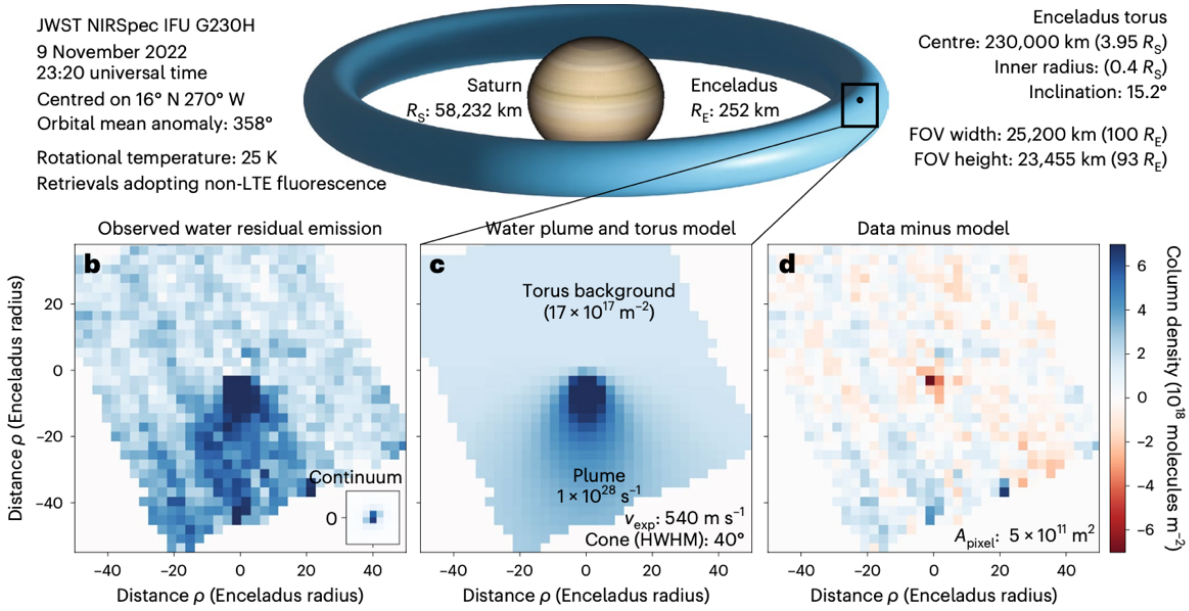


Figure 3. James Webb Space Telescope maps of the water eruptions from the south pole of moon Enceladus. Enceladus is smaller than one pixel in these maps, but the plume extends over a much larger region (from Villanueva et al. 2023).

In this project, students will analyze these brand-new JWST observations of Enceladus. The main goal of the project is to determine whether the plume activity has changed compared with earlier JWST observations obtained in 2022 and 2024.

The project will include the following steps:

1. Download JWST NIRSpec data from the online archive and become familiar with the available data products.
2. Adapt existing spectral-analysis code to analyze the signal containing information about the water-vapor plume emissions.
3. Account for the observing conditions of the new data in the analysis. This includes the geometry of the planet–moon system and the intensity of the solar flux on the observation date.
4. Produce gas-emission maps similar to the one shown in the figure above and compare them with previously derived JWST maps.

By carrying out this project, students will contribute to ongoing space-science research while gaining hands-on experience in the analysis of space-telescope data and scientific data-analysis methods.

Project J3: Schur decomposition of Hall-magnetohydrodynamic numerical simulations of turbulent space plasmas.

Supervisor: Luca Sorriso-Valvo, lucsv@kth.se, Electromagnetics and Plasma Physics

Spacecraft measurements have shown that solar wind and planetary magnetospheres are continuously heated through dissipation of kinetic and magnetic energy. Unlike common fluids and plasmas, space plasmas are extremely rarefied and hot, so that collisions among particles are nearly negligible. Hence, dissipation is not produced by viscous-resistive effects. Rather, complex field-particle interactions, such as kinetic plasma instabilities, magnetic reconnection or wave-particle resonances, result in energy conversion from ordered forms to heating. Understanding the nature of the processes that dissipate the enormous amount of energy injected by solar, coronal, and large-scale solar wind structures, and heat the solar wind plasma, is one of the major challenges of space physics. One crucial open question is whether the properties of the local energetic small-scale structures generated by turbulence activate specific dissipation mechanisms. Using the Magnetospheric Multi-Spacecraft (MMS) measurements in the highly turbulent terrestrial magnetosheath, possible relations between turbulent structures and dissipation were observed. Motivated by these discoveries, we aim to obtain detailed geometrical characteristics of the turbulent fluctuations and explore their potential effects on energy dissipation. Finding such effects is a crucial step to unveil the underlying kinetic processes and their activation conditions, providing closure to the outstanding question of solar wind heating.

Recent studies used Schur decomposition of turbulent velocity and magnetic gradients to separate regions of space dominated by shear, strain, and rotation. This allows for strong geometric characterization of dissipating structures. Those studies focused on a MHD plasma description. However, for space plasmas, the Hall-MHD framework represents a more suitable bridge between the fully fluid (MHD) approximation and the more complete kinetic description of the dynamics, including effects of decoupling of ions and electrons. Here, we address the possibility of using Schur decomposition on Hall-MHD fields and explore its ability to identify dissipation regions.

This project will thus address the following questions:

1. How does the Schur decomposition perform for Hall-MHD simulations?
2. Do different geometrical properties of fluctuations correspond to different dissipation levels?

The student will learn basic concepts of space plasma turbulence and acquire competences in obtaining, managing and analyzing numerical simulation data, performing the Schur decomposition, and computing turbulent energy transfer and dissipation. State-of-the-art, high-resolution Hall-MHD numerical simulations of turbulent space plasmas under different regimes will be used to perform a characterization of the dissipative structures according to their geometrical properties. This will be compared with energy transfer and dissipation estimates. Optimally, machine-learning clustering techniques will help determine the existence of correlations between these parameters. The analysis will be performed using a programming language of choice, or an existing suite of analysis tools. The work will be the continuation of ongoing research by KTH scientists.

Projekt J4: Observing pulsating aurora with a ground-based GNSS network

Supervisor: Mykola Ivchenko, nickolay@kth.se , Electromagnetics and Plasma Physics



Pulsating aurora is an exciting phenomenon during which the brightness of auroral displays changes intensity quasiperiodically, turning “on” and “off” at periods of few seconds to few tens of seconds. The source of pulsating aurora is thought to be magnetosphere, where trapped electrons interact with plasma waves and get diverted to the atmosphere, producing optical emissions. Apart from optical emissions, electron produce ionization in the ionosphere. Enhanced ionization can be detected by Global Navigation Satellite System (GNSS) reference stations. Electron concentration integrated along the line of sight to the satellite, known as Total Electron Content (TEC), can be obtained from the analysis of dual frequency GNSS data. The advantage of using GNSS TEC observations is the broad spatial coverage and independence of the clouds (which otherwise limit optical imaging). In this project you will work with GNSS data obtained with the Swedish SWEPOS network (https://www.lantmateriet.se/sv/geodata/gps-geodesi-och-swepos/swepos/) to analyze some pulsating aurora events.

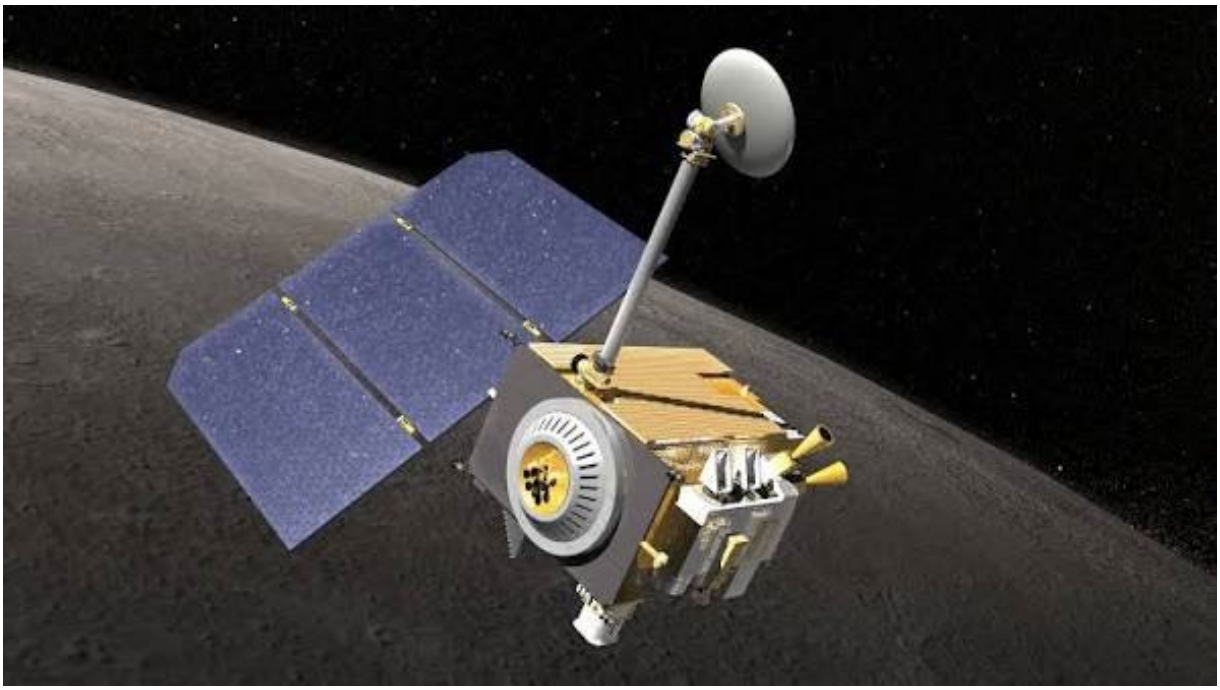
The tasks in this project include:

- Getting familiar with the GNSS receiver data analysis to obtain TEC variations
- Getting familiar with the pulsating aurora phenomenon
- Identifying pulsating aurora events for further studies
- Analysing the events using a large number of SWEPOS stations to characterize the spatial and temporal properties of the pulsating aurora
- Comparing the GNSS TEC variations with optical imaging if relevant

Project J5: Lunar orbit propagation and error analysis

Supervisor: Guillermo Peña Martínez, gupm@kth.se , Electromagnetics and Plasma Physics

The accuracy of a lunar orbit propagator depends strongly on the orbit being flown: the Moon's irregular gravity field (mascons) and Earth's third body pull make some orbits evolve, and accumulate error, much faster than others. In this project, the students will test the orbit propagator developed for a PhD project on a set of representative orbits flown by lunar spacecraft, from low circular to highly elliptical ones. Each orbit is propagated and compared with the spacecraft's reconstructed trajectory from the JPL database (e.g. LRO and GRAIL). The students will then change the starting point of each orbit (initial position along the orbit) to show how the orbital evolution and the error growth depend on the initial conditions. The result is a clear picture of where the propagator performs well and which orbits and starting conditions are the most demanding.



The tasks in this project include:

- Getting familiar with lunar orbital dynamics, main perturbations, and reference frames.
- Getting familiar with orbital propagation.
- Propagating each orbit and analyzing the evolution of its orbital elements.
- Error decomposition and computation (RTN frame)
- Analysis of the errors: how each component grows over time and how it differs between orbit types.

Project J6: Exploring Venus' upper atmosphere with Pioneer Venus Orbiter Data

Supervisor: Shane Carberry Mogan, shanecm@kth.se, Electromagnetics and Plasma Physics

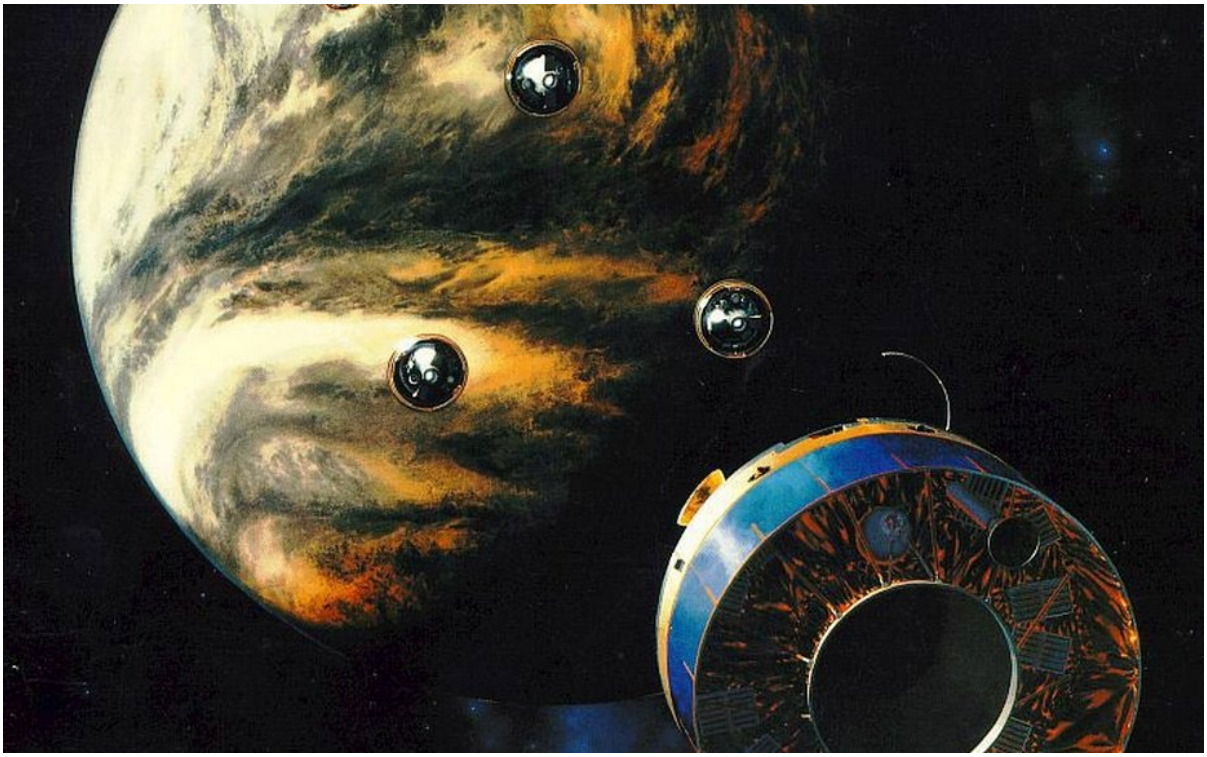


Figure 1. Artist's illustration of NASA's Pioneer Venus mission at Venus, showing the Pioneer Venus Orbiter in the foreground and the atmospheric entry probes approaching the planet.

Venus has a dense carbon-dioxide (CO_2) atmosphere, but at high altitudes the gas becomes increasingly rarefied and strongly influenced by solar radiation, atmospheric dynamics, and interaction with the solar wind. Because Venus lacks a global intrinsic magnetic field, its upper atmosphere and ionosphere are directly exposed to the surrounding space environment. Understanding how the thermosphere, exosphere, and ionosphere vary is therefore important for understanding the long-term evolution of the planet.

NASA's Pioneer Venus Orbiter (PVO) studied Venus from 1978 to 1992. The spacecraft's Orbiter Neutral Mass Spectrometer (ONMS) measured the densities of important neutral species including CO_2, CO, N_2, N, O, and He, while the Orbiter Ion Mass Spectrometer (OIMS) measured the composition and density of various ions in the ionosphere. These historic datasets provide vertical density profiles under different solar zenith angles, local solar times, and solar-activity conditions, making them well suited for a modern data-analysis project.

In this project, you will use PVO-ONMS and OIMS data to establish a statistical baseline description of Venus's upper atmosphere and ionosphere, examining how their structure varies with parameters such as altitude, local solar time, solar zenith angle, year, and solar activity. You will then search for departures from this baseline that may reveal additional atmospheric or plasma processes.

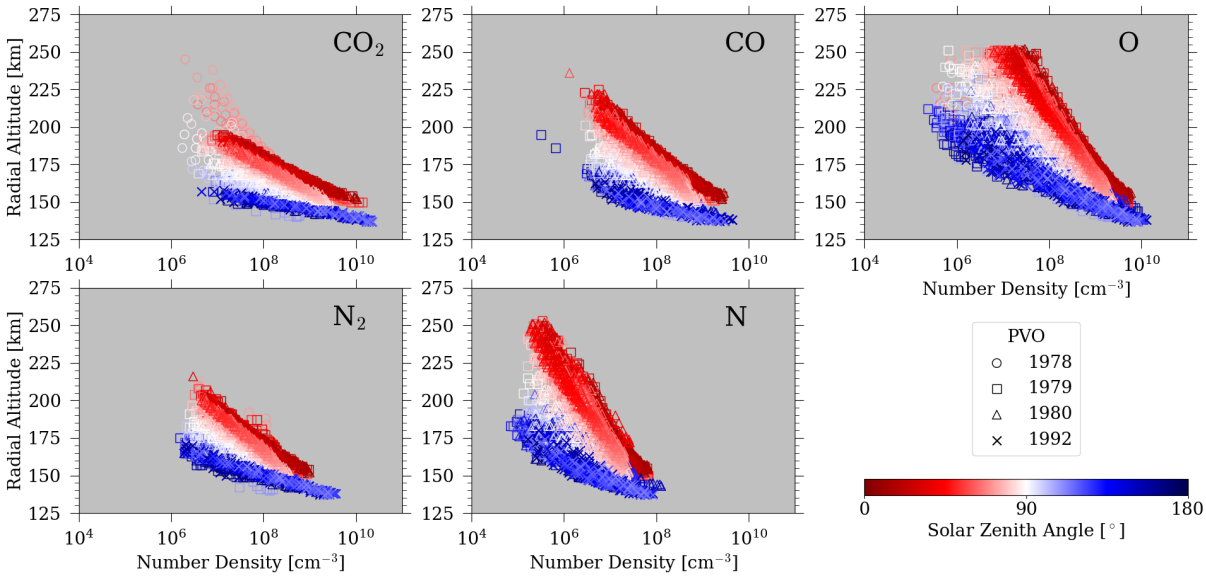


Figure 2. Example PVO/ONMS neutral-density profiles for CO_2, CO, O, N_2, and N. Symbols indicate observing year and color indicates solar zenith angle. The large orbit-to-orbit variability provides a natural starting point for separating the mean atmospheric structure from dynamical and externally driven variations.

The tasks include:

1. **Processing ONMS and OIMS measurements together with spacecraft and observing geometry**, including altitude, latitude, local solar time, and solar zenith angle, across the available PVO periapsis passes.
2. **Constructing representative vertical density profiles** for the major neutral and ion species and determining how the mean atmospheric and ionospheric structure changes with observing conditions, for example between the dayside and nightside or between periods of different solar activity.
3. **Fitting smooth background profiles and examining the residuals** for coherent structures such as density enhancements, depletions, or oscillations. These structures can then be investigated for signatures consistent with atmospheric dynamics such as gravity waves, or with departures from a simple thermal or hydrostatic population, for example enhanced high-altitude densities.
4. **Comparing the observed variability with solar-activity proxies** and investigating whether changes in the neutral atmosphere are accompanied by corresponding changes in the ionosphere.

By carrying out this project, you will gain experience in working with archival spacecraft data, atmospheric-profile analysis, statistical methods, curve fitting, and the interpretation of planetary-atmosphere and ionosphere measurements.

Context K:
Networked Artificial Intelligence



Over the past decades, we have experienced a series of revolutions in computing, information, and communication technologies—starting with the invention of computers, followed by networking, and then wireless communication. Now, we stand on the brink of the next great leap: the "all-connected and digitalized world," where networks of interconnected devices will drive automatic data analysis and decision-making, transforming everyday objects into intelligent systems.

By 2030, it is predicted that the number of connected devices will reach 100 billion, ushering in a multi-trillion-dollar economy. One defining feature of this fourth revolution is the massive volume of data being generated, particularly by Internet of Things (IoT) devices, smartphones, and social networks. In fact, 90% of the world's data has been produced in just the last few years, creating an urgent need for advanced data analysis methods like **Machine Learning (ML)** and **Artificial Intelligence (AI)**.

While ML has already achieved remarkable feats—such as Google's AlphaGo, which trained on 30 million possible moves to defeat a Go grandmaster—these accomplishments rely heavily on centralized data processing in high-powered environments like data centers. However, the reality of the fourth revolution is different: data will be distributed across millions of nodes, from smartphones to sensors, each with limited computing power and bandwidth. Traditional AI/ML methods, designed for centralized systems, struggle in this new landscape, particularly in networks like IoT or 5G, where bandwidth is scarce, and communication protocols are diverse.

This is where **Networked AI** comes into play. The challenge is to design AI/ML algorithms that can function efficiently across distributed systems. Instead of relying on central servers, these algorithms need to operate collaboratively across many devices, each making decisions based on its local data and only exchanging essential information with others. This requires addressing key technical hurdles, such as communication delays, bandwidth constraints, and the need for coordination between nodes with different computational capacities.

One major challenge is the fundamental bandwidth limitation. As more devices transmit data, congestion becomes a significant concern in a shared wireless medium. Emerging low-latency communication technologies rely on short packets with limited data capacity, further complicating ML over networks. IoT systems, such as smart grids or smart cities, will often operate in environments with unreliable links, low data rates, and transmission delays—whether it is underwater IoT sensors or communication networks within the human body, which transmit only a few bits per second.

This bachelor thesis context offers you the opportunity to work at the forefront of this exciting field. You will explore cutting-edge techniques for **Networked AI**, studying the latest advancements in distributed optimization and investigating how to train ML models across a network of devices. Nodes will collaborate by learning from their local data and sharing only essential information with others using wireless communications. Using the latest tools, simulations, and theoretical models, you will develop solutions that could shape the future of intelligent, distributed networks.

If you are passionate about AI, data science, and networking, this is your chance to be part of a transformative technological revolution through one of the following exciting projects:

- Understanding Wireless Networks and Machine Learning (ML) to Develop Network-Assisted ML Solutions
- Understanding Key 6G Technologies and Their Integration for a Unified Framework
- Distributed Image Classification with Federated Learning
- First-Order Optimization Methods for Machine Learning
- Online Convex Optimization
- Digital Modulation for over the air computations
- Semantic-Aware GNNs for Networked Vehicles
- Conformal Prediction for Reliable Channel Estimation
- Distributed Optimization in Machine Learning

We describe the projects in detail on the following pages.

Project K1: Understanding Wireless Networks and Machine Learning (ML) to Develop Network-Assisted ML Solutions

Supervisor: Swapnil Sadashiv Shinde (ssshinde@kth.se), Carlo Fischione (carlofi@kth.se)

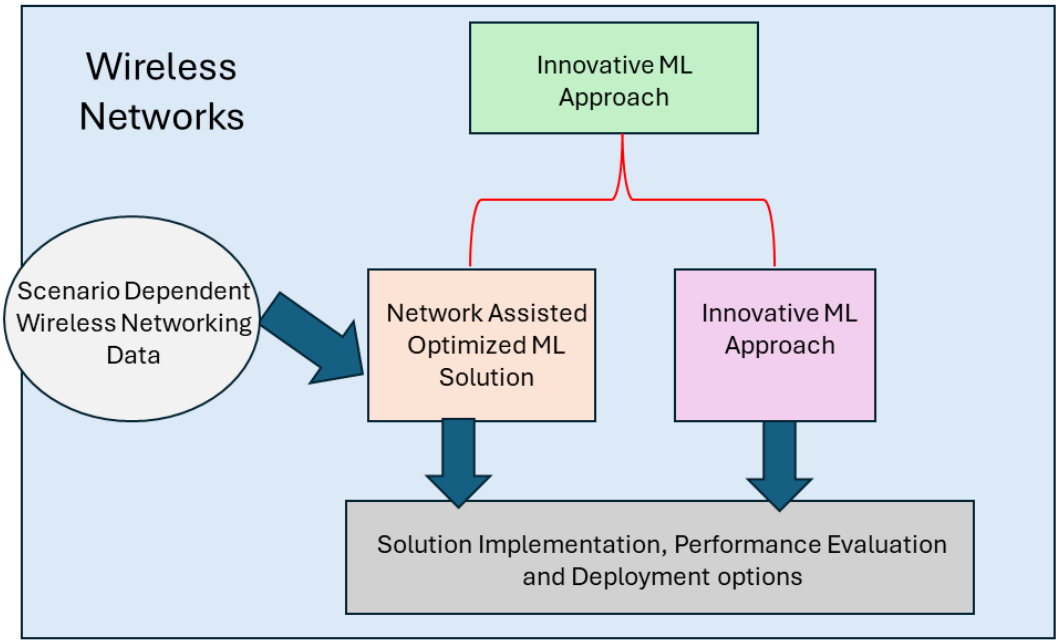


Figure 4: Network Assisted Innovative ML Solution

Over the past decade, Machine Learning (ML) has gained significant attention in the wireless research community, especially for tackling challenging networking problems such as resource allocation, user scheduling, and mobility management. However, when designing ML solutions, it is crucial to consider the resource-constrained nature of wireless networks along with strict requirements on latency, reliability, and scalability. These limitations often become bottlenecks for deploying ML-based solutions in real-world wireless systems.

A promising direction to address these challenges is the development of **network-assisted ML solutions**, which incorporate network resources, user demands, and system capabilities into the learning process while solving key networking problems.
In this project, the student will learn the fundamentals of **wireless networks** and **machine learning** necessary to understand network-assisted ML solutions.

By the end of the project, the student will be able to:

1. Understand the basics of wireless networks and their limitations.
2. Learn fundamental ML concepts relevant to wireless networking.
3. Study and analyze existing ML and network-assisted ML solutions.
4. Summarize the concepts through reports or presentations, including basic performance insights.

Project K2: Understanding Key 6G Technologies and Their Integration for a Unified Framework

Supervisors: Swapnil Sadashiv Shinde (ssshinde@kth.se), Carlo Fischione (carlofi@kth.se)

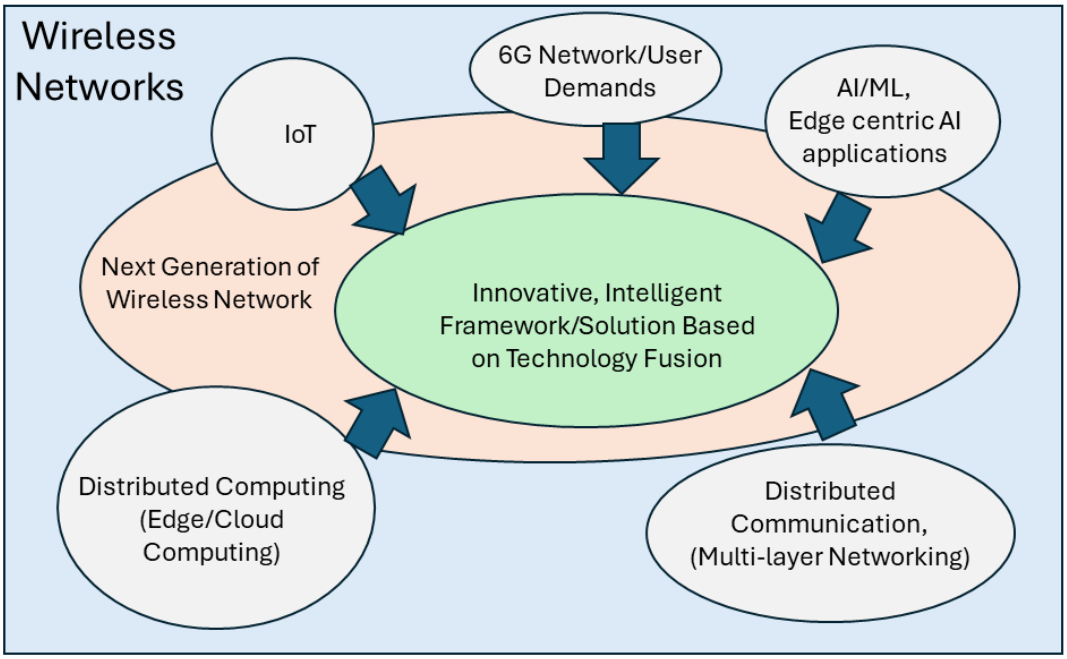


Figure 5: Unified 6G Framework Integrating Emerging Technologies

The upcoming generation of wireless networks, **6G**, is expected to enable a fully connected and intelligent world serving the needs of users in the 2030s. To achieve this vision, 6G will need to support novel services and applications with embedded intelligence.
Several technological trends are expected to shape 6G, including:

- **Internet of Things (IoT)**
- **Multi-layer networks with heterogeneous nodes**
- **Edge-centric distributed AI**
- **Distributed computing**

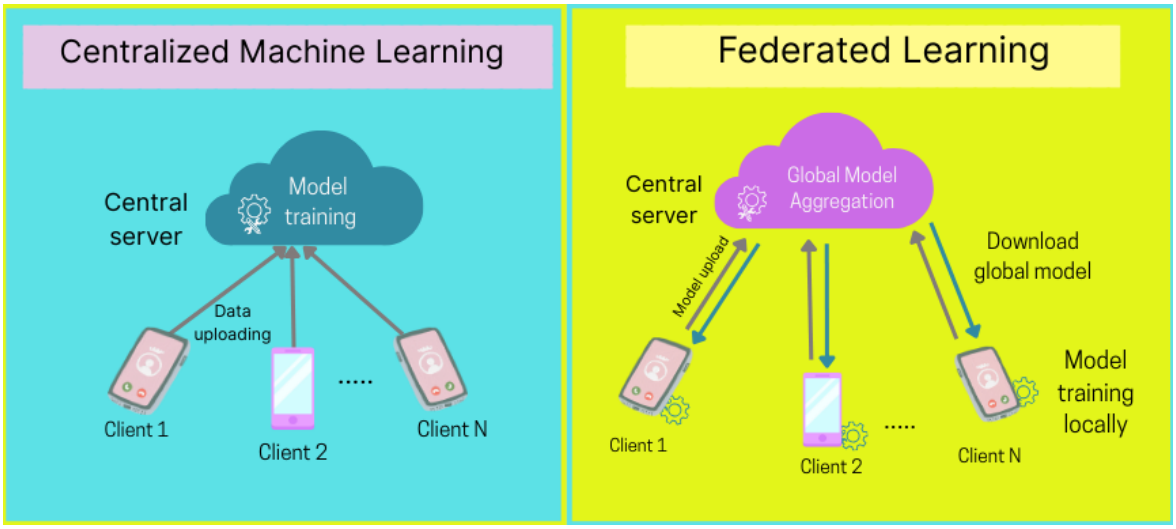
Understanding the requirements and interactions of these technologies is essential for designing a **unified 6G framework** that ensures their efficient coexistence and integration.
In this project, the student will focus on **learning, analyzing, and summarizing existing knowledge** about 6G technologies and investigating the potential for designing a framework that enables their efficient coexistence.

By the end of the project, the student will be able to:

1. Understand the fundamentals of key 6G technologies and their roles in future networks.
2. Learn the technical requirements, strengths, and limitations of each technology.
3. Analyze potential dependencies, conflicts, and complementarities among these technologies.
4. Summarize findings in a report, including recommendations for integration and directions for future research.

Project K3: Distributed Image Classification with Federated Learning

Supervisors: Milica Jankov (milicaj@kth.se), Carlo Fischione (carlofi@kth.se)



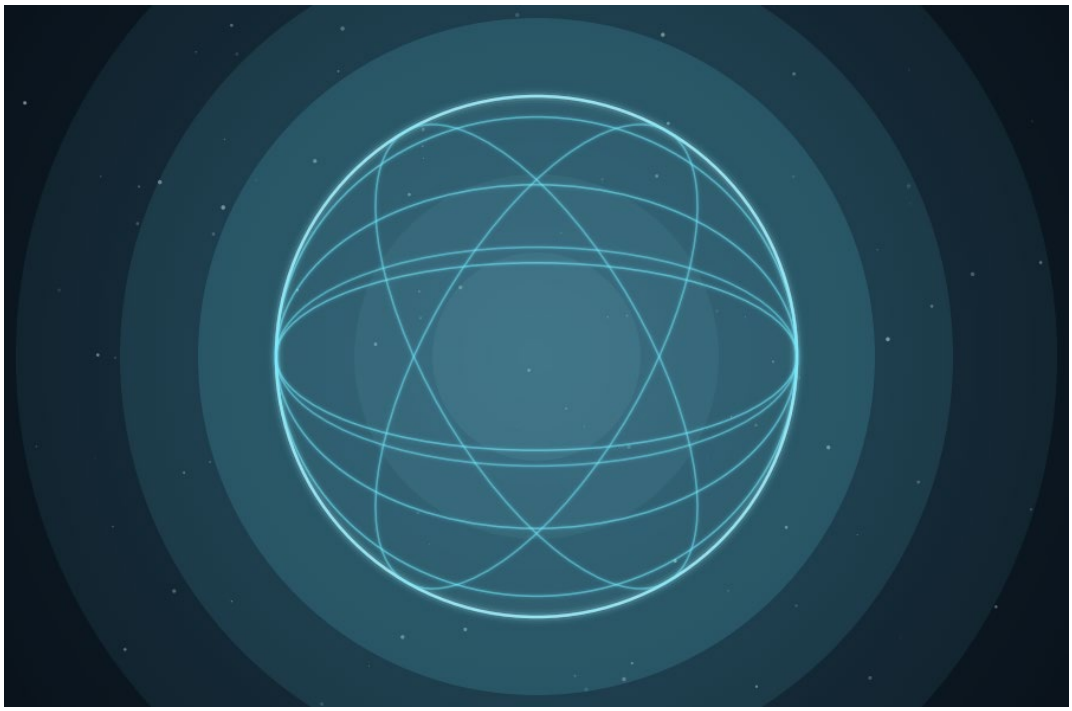
The rapid increase in connected devices, including smartphones and IoT sensors, has generated massive volumes of data that are inherently distributed across networks. Traditionally, machine learning models are trained in centralized environments, requiring all data to be aggregated at a single location. However, in many real-world scenarios, this approach is inefficient or even infeasible due to privacy concerns, communication costs, and limited bandwidth. Federated learning addresses these challenges by enabling devices to collaboratively train a shared model without transmitting their raw data to a central server.

This project aims to introduce the fundamentals of federated learning through a simple image classification task. By simulating multiple devices, each using only a subset of the dataset, students will explore how federated learning works in practice and how it compares to a centralized setting. Specifically, the main objectives of this study are as follows:

- Gain a basic understanding of federated learning and its motivation.
- Train a simple neural network on the widely used MNIST image dataset using simulated distributed clients.
- Compare the centralized approach (all data collected at a single location) with the federated approach (data distributed across clients).
- Evaluate how the number of participating clients and the distribution of data among them affect the final model accuracy.

Project K4: Learning World Models for Distributed and Networked Systems

Supervisors: Simone Bortolotti (simonebo@kth.se), György Dán (gyuri@kth.se)
Department of Network and Systems Engineering



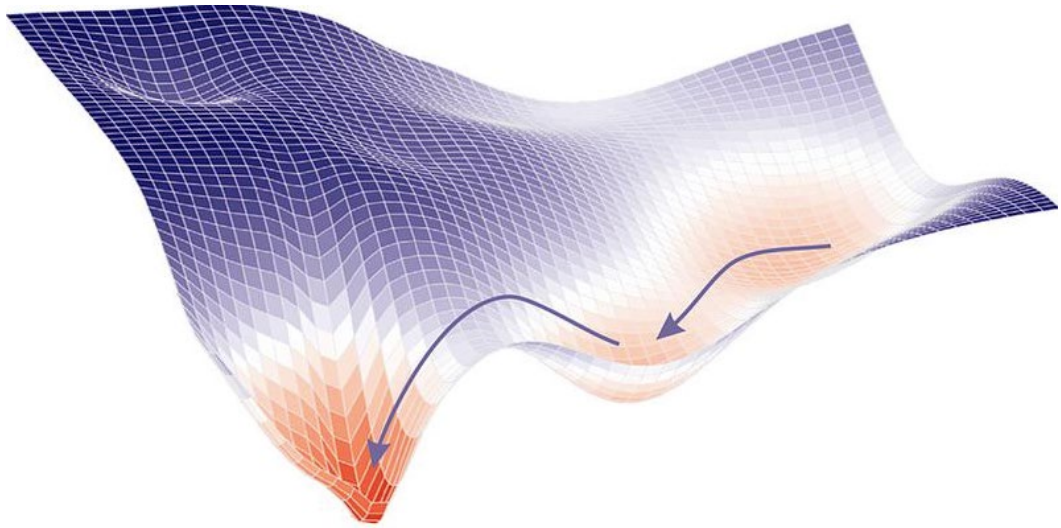
A world model is a learned internal representation of how an environment evolves over time. A system equipped with a world model can predict future states, simulate outcomes of different actions, and plan ahead, without repeatedly interacting with the real environment. World models have become an important approach in model-based reinforcement learning and are increasingly being explored for robotics and sample-efficient control of complex systems.
Learning a world model becomes more challenging when observations of the environment are distributed across many devices: a wireless channel, a fleet of vehicles, a sensor network, or a smart grid, where observations are generated locally and no single node sees the full picture. Training a world model in this setting requires deciding what to observe, what to communicate, and how often, since devices are limited in computation, and links are constrained in bandwidth, delay, or reliability. This project explores how world models can be learned and kept accurate in such distributed, communication-constrained settings, where devices share only the information necessary to maintain a useful model.

By the end of the project, the student will be able to:

- Understand the fundamentals of world models, how they are learned and how they support prediction, simulation, and planning.
- Study how world models can be trained or updated in distributed settings, where devices have limited computation and communication resources.
- Implement a small-scale demonstrator, e.g., a world model of a simple dynamical system, wireless channel, or IoT process, and evaluate its predictive accuracy.
- Analyze the trade-off between communication cost and model accuracy, summarizing the findings in a report supported by simulations and plots.

Project K5: First-Order Optimization Methods for Machine Learning

Supervisors: Seyedamirreza Kazemi (seykaz@kth.se), Carlo Fischione (carlofi@kth.se)



First-order optimization methods, such as Gradient Descent and its variants, are central to modern **machine learning** and **numerical optimization**. They form the backbone of training algorithms for models ranging from logistic regression to deep neural networks. By relying on gradient information to minimize loss functions, these methods make it possible to fit models to data efficiently and at scale. A solid understanding of their behavior is essential for both the theoretical foundations of machine learning and for practical applications.
This proposal outlines a bachelor thesis project focused on exploring the **basic principles, practical implementations, and comparative performance of first-order optimization algorithms**.

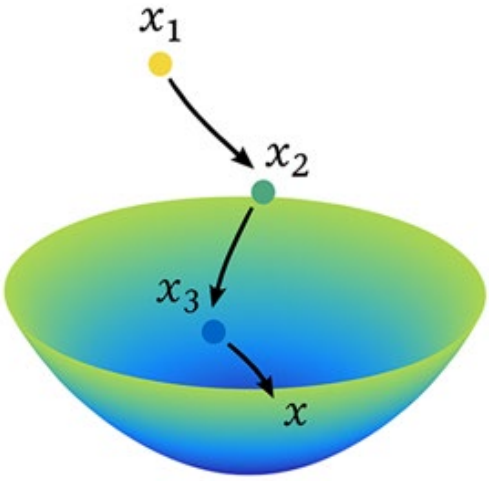
The primary objectives are:

- **Implement Gradient Descent and several of its variants** (e.g., Momentum, Nesterov, SGD, Adam) from scratch.
- **Apply these methods to benchmark problems**, e.g., Logistic Regression on a dataset and the Rosenbrock function.
- **Compare the performance** of the different methods in terms of convergence speed, robustness, and sensitivity to hyperparameters.

By the end of the project, students should have a clear understanding of how different first-order optimization methods perform in practice, along with hands-on experience implementing them and comparing their strengths and weaknesses. The main deliverables will be a concise report with explanations, plots of the results, and well-structured Python or MATLAB code.

Project K6: Online Convex Optimization

Supervisors: Krishnendu S. Tharakan (tharakan@kth.se), Carlo Fischione (carlofi@kth.se)



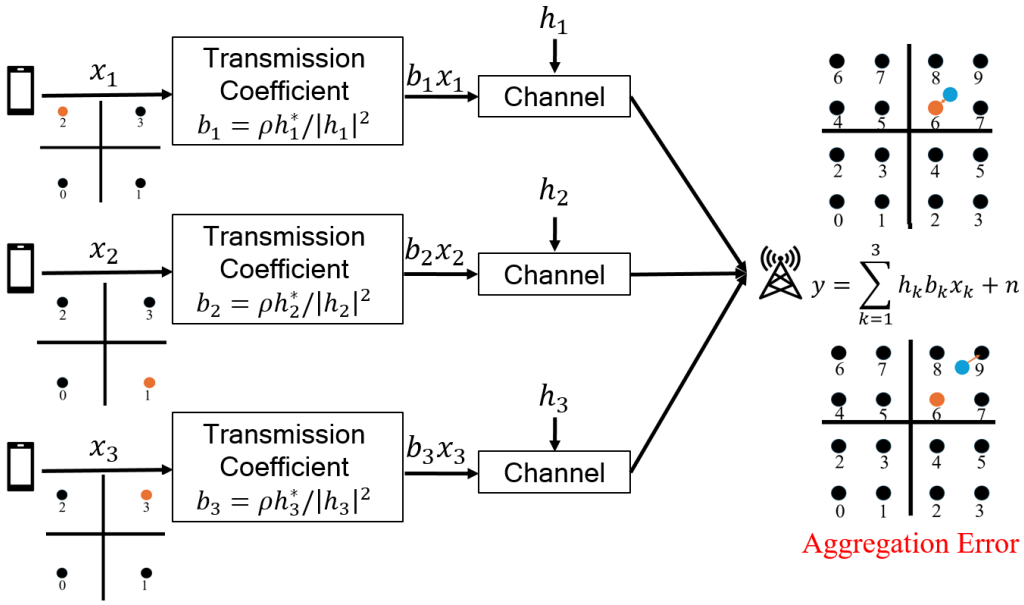
Many machine learning and resource allocation problems can be framed as sequential decision problems under uncertainty. At each step, an agent must make a decision without knowing the future, and then observe a loss or reward. Examples include portfolio optimization, online recommendations, and dynamic resource allocation. Online Convex Optimization (OCO) provides a unifying framework for such problems: decisions are chosen from a convex set, losses are convex functions revealed sequentially, and the quality of an algorithm is measured by its regret compared to the best fixed decision in hindsight.

This project explores the online learning version of the knapsack problem with switching costs. This project explores fundamental algorithms in OCO, such as Online Gradient Descent (OGD) and Follow-The-Regularized-Leader (FTRL), and evaluates their theoretical guarantees and practical performance.

- Introduce the online convex optimization framework and regret analysis.
- Implement core OCO algorithms, such as OGD and FTRL.
- Prove basic regret bounds for OGD (e.g., $O(\sqrt{T})$).
- Simulate and compare algorithms under different demand sequences.

Project K7: Digital Modulation for over the air computations

Supervisors: Zeyang Li (zeyangl@kth.se), Carlo Fischione (carlofi@kth.se)



In traditional wireless systems, users take turns (use separate time/frequency bands) to transmit so their signals do not overlap (orthogonal communication).
Over-the-Air Computation (AirComp) is different: several devices transmit at the same time and the wireless channel naturally adds their signals. If we design the transmissions properly, the receiver can directly estimate an aggregate (e.g., sum or average) of the users' data. This is useful for sensor networks, Internet of Things monitoring, and federated learning where only the aggregate is needed.

Problem statement

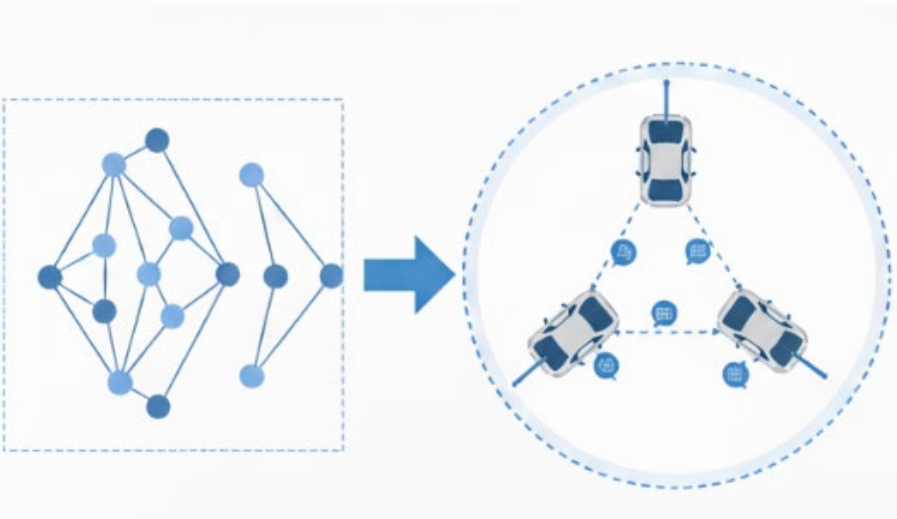
We want to find which digital modulation works better for Digital AirComp. You will compare basic digital modulation (e.g., BPSK, 4-QAM, optionally 16-QAM) and measure the aggregation accuracy (e.g., mean-squared error, MSE) under different signal-to-noise ratio (SNR) and number of users.

What you will learn

- Basic digital modulation: BPSK and 4-QAM (Quadrature Amplitude Modulation).
- Conceptual contrast: orthogonal communication vs AirComp.
- Simple performance drivers: SNR, number of users; how they influence aggregation accuracy.
- build a small simulator (MATLAB or Python), generate plots, and explain results.

Project K8: Semantic-Aware GNNs for Networked Vehicles

Supervisors: Minyi Wan (minyiw@kth.se), Carlo Fischione (carlofi@kth.se)



Graph-structured learning naturally models connected vehicles and road infrastructure: vehicles as nodes, V2X links as edges. On the other hand, the latest development in semantic communication brings the opportunity to boost the development from Classical Shannon's Theory. This project offers an introduction to GNNs and semantic communication while adding a small, vehicular-flavored graph for semantic communication.

The main objectives are as follows:

- Literature Review:
  - Conduct a literature review of GNNs algorithms, including basic GNN, GCN, GAT, etc.
  - Study general ideas about semantic communication concepts, understand the difference with three levels from Shannon-Weaver Explanation [1].
- Model Construction
  - Build and simulate a simple GNN (chose from literature review) with simulated noise in a vehicle network.
  - Build benchmark models (e.g. CAE, CNN, etc.) [2] for the comparison.
- Result Evaluation
  - Analyze how each part of GNN contributes to the predicted results.
  - Carry out comparison experiments with GNNs and other methods.

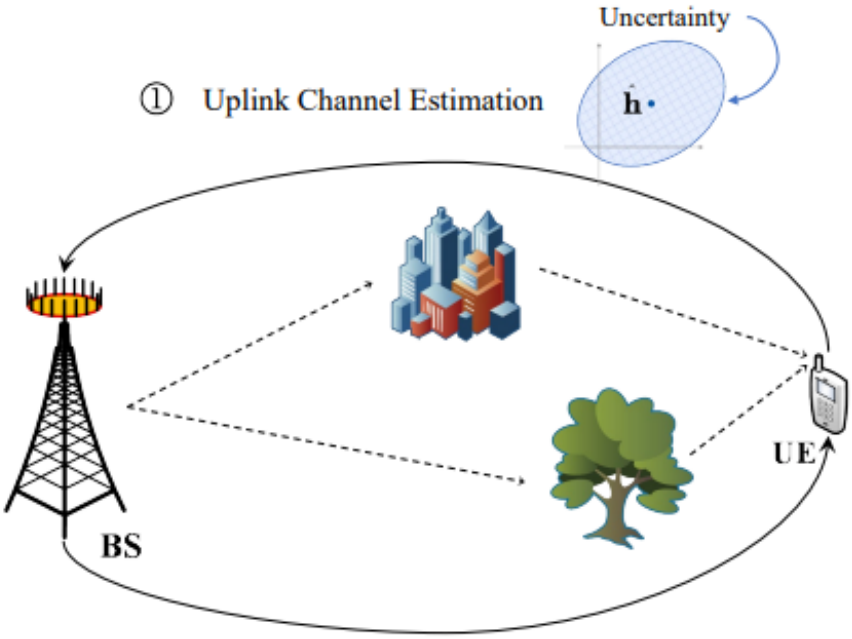
References:

[1] W. Weaver, "Recent Contributions to The Mathematical Theory of Communication," in *The Mathematical Theory of Communication*, C. E. Shannon and W. Weaver, Urbana, IL, USA: Univ. of Illinois Press, 1949, pp. 1–28.

[2] E. Eldeeb, M. Shehab, and H. Alves, "A Multi-Task Oriented Semantic Communication Framework for Autonomous Vehicles," 2024, *arXiv:2403.12997*. [Online]. Available: https://arxiv.org/abs/2403.12997

Project K9: Conformal Prediction for Reliable Channel Estimation

Supervisors: Xin Su (xisu@kth.se), Carlo Fischione (carlofi@kth.se)



Channel estimation is a key step in wireless communication systems, but the estimates are affected by noise and uncertainty. To make reliable decisions, it is important not only to provide a point estimate of the channel but also to quantify the uncertainty of the estimate. Conformal Prediction (CP) is a general statistical method that can wrap around any estimator to provide prediction sets with guaranteed reliability.

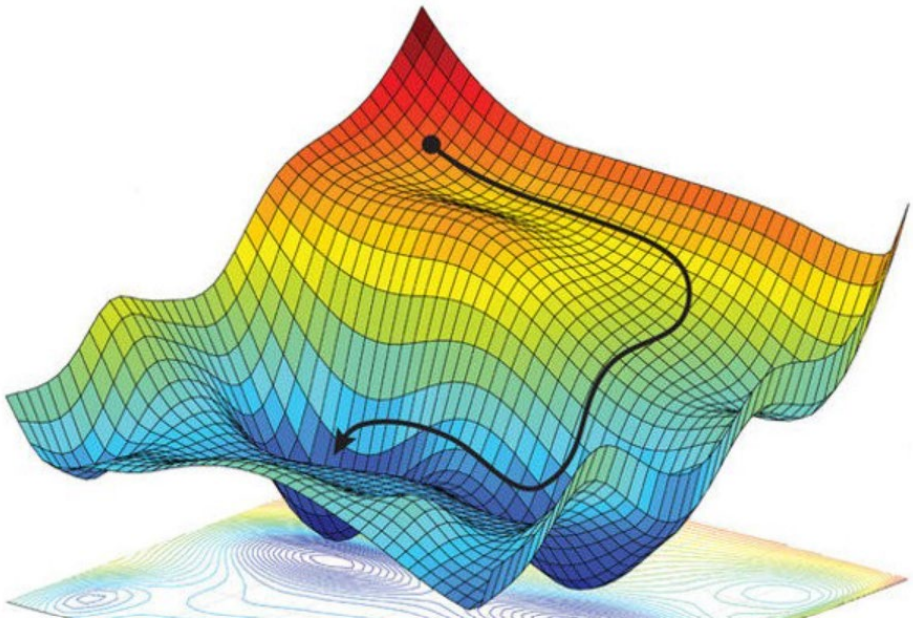
This project introduces CP and applies it to channel estimation problems in wireless systems. The aim is to study whether CP can achieve the desired reliability levels under different noise conditions.

The main objectives are as follows:

- Study the basic principles of CP from tutorials and introductory papers.
- Generate synthetic wireless datasets with different types of noise, such as additive white Gaussian noise and heteroscedastic noise.
- Implement least squares (LS) and multilayer perceptron (MLP) channel estimator optimized using MSE loss.
- Apply CP to these estimators and evaluate whether the resulting prediction sets achieve the required coverage levels.

Project K10: Distributed Optimization in Machine Learning

Supervisors: Seif Hussein (seifh@kth.se), Carlo Fischione (carlofi@kth.se)



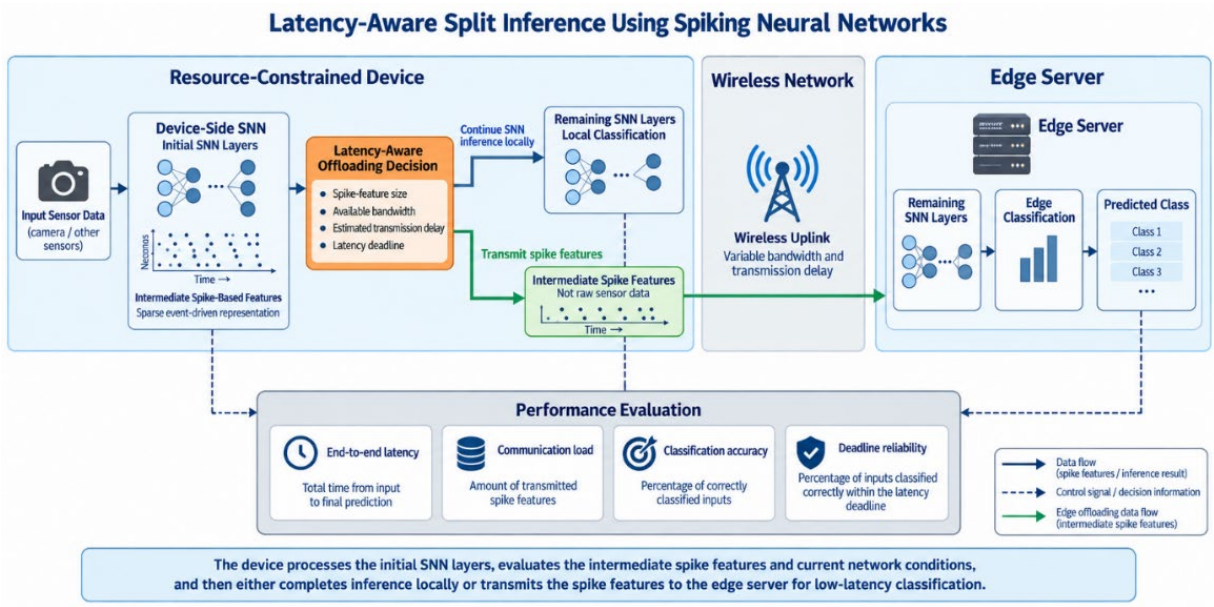
In the era of big data, the processing of high-dimensional datasets has become practically infeasible for traditional centralized computing systems. Machine learning models require immense computational resources to process and analyse this data, leading to significant challenges in scalability and efficiency. Distributed and decentralized optimization algorithms have thus emerged to address these challenges, enabling the processing of large-scale data across multiple computing nodes or networks.

This project aims to explore various distributed and decentralized optimization algorithms proposed in the literature. In particular, the expected outcomes of the project are the following:

- Analyze optimization algorithms commonly used in various machine learning settings
- A categorization of some state-of-the-art distributed/decentralized optimization algorithms based on mathematical properties and application areas
- Simulations to highlight differences between algorithms

Project K11: Latency-Aware Edge Offloading with Spiking Neural Networks

Supervisors: Shraddha Tripathi (shrtri@kth.se), György Dán (gyuri@kth.se), Department of Network and Systems Engineering



Resource-constrained devices increasingly perform perception tasks using cameras and other sensors. Sending all captured data to an edge server increases communication load and latency, whereas fully local processing may exceed the device's computational and energy capacity. Adaptive edge offloading addresses this trade-off by selecting where inference should be completed. Spiking Neural Networks (SNNs) process information through sparse, event-driven spikes. In this project, the initial SNN layers will run on the device to generate intermediate spike features. Based on their size and current network conditions, the device will either complete inference locally or transmit the features to an edge server for classification.
The objective is to determine whether SNN-based split inference can reduce communication load and end-to-end latency while maintaining classification accuracy. A simple rule based on feature size, available bandwidth, and estimated transmission delay will control the offloading decision. Reliability will be measured as the percentage of inputs classified correctly within a latency deadline. The project will involve:

- Implementing a small SNN using an existing framework such as snnTorch.
- Dividing the model at one selected intermediate layer.
- Implementing local inference and spike-feature offloading.
- Developing a simple latency-aware offloading rule.
- Comparing local inference, raw-input offloading, and spike-feature offloading.
- Evaluating accuracy, transmitted data volume, latency, and deadline reliability.

Related Work

1. C. Feng et al., "Computation Offloading in Mobile Edge Computing Networks: A Survey," *Journal of Network and Computer Applications*, vol. 202, 2022.
2. Y. Liu, Z. Qin, and G. Y. Li, "Energy-Efficient Distributed Spiking Neural Network for Wireless Edge Intelligence," *IEEE Transactions on Wireless Communications*, vol. 23, no. 9, 2024.

Context L: Cybersecurity



Cybersecurity includes a collection of methods designed to protect systems, networks, and services from external threats. Businesses and organizations employ cybersecurity professionals to protect their confidential information, maintain employee productivity, and enhance confidence in products and services.

Key properties of cybersecurity are privacy, integrity, and availability. Privacy means data can be accessed only by authorized parties; integrity means information can be added, altered, or removed only by authorized users; and availability means systems, functions, and data must be available on-demand according to agreed-upon parameters. An important element of cybersecurity is the use of authentication mechanisms, which allow to securely identify users or processes.

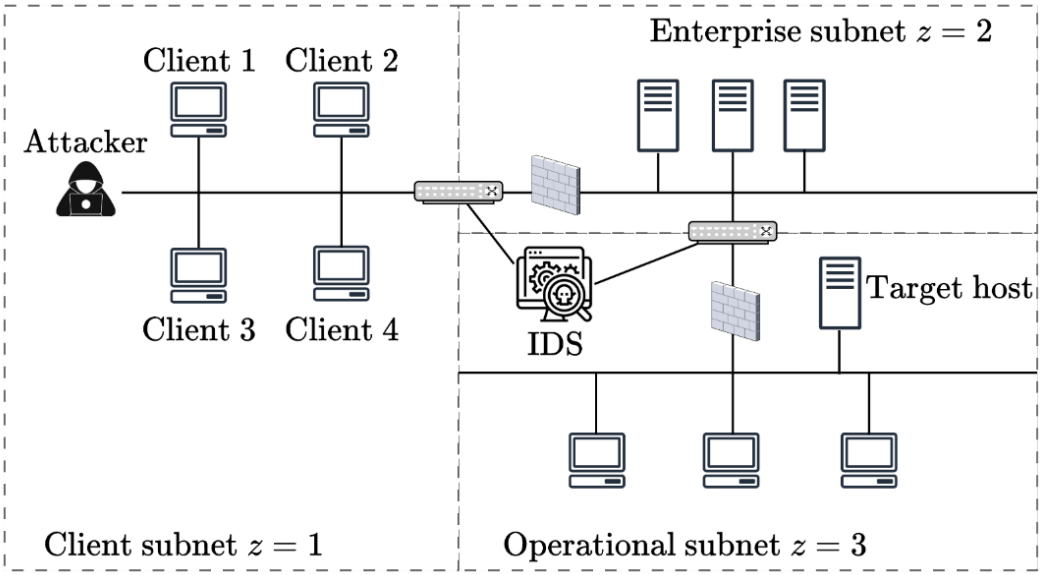
Well-known attacks that compromise the protect systems, networks, and services of an organization include Denial of Service (DOS), installation of malware, man-in-the-middle attack, and phishing. Other types of cyberattacks include cross-site scripting attacks, password attacks, eavesdropping attacks (which can also be physical), SQL-injection attacks, and birthday attacks.

From an engineering point of view, a variety of scientific methods can be used to protect a system or to identify and repel an attack. These include cryptography, formal methods, statistical techniques, and AI. We offer projects where students develop and evaluate state-of-the art approaches to securing systems and preventing attacks:

1. Automated Intrusion Response for IT systems
2. Designing and Simulating Cyber Attacks on Drones
3. Adversarial Attacks on Multi-Agent Reinforcement Learning System
4. Robustness of RNN-Based World Models
5. Adversarial Attack and Defense for Deep Vision Models
6. Formal Verification in Cybersecurity
7. Adding Static Information-Flow Checks to Secure AI Agents
8. Side Channel Analysis on Frodo Key Encapsulation Mechanism

Project L1: Automated Intrusion Response for IT systems

Supervisors: Rolf Stadler (stadler@kth.se), Duc Huy Le (dhle@kth.se)



Due to the rapid evolution of digital technologies and the growing sophistication of cyber threats, automated cyber defense systems become essential for detecting and responding to attacks in real time. This project focuses on an intrusion response use case in which the defender must learn effective strategies to achieve two key objectives: protecting the system against Advanced Persistent Threats (APTs) and maintaining service availability for external clients.

The main tasks of the project are to formulate a scenario for a small enterprise, including network and system configuration, attacker behavior, alerts and observations available to the defender, as well as response actions the defender can take. Second, the scenario will be modeled and the defender problem formulated using using the statistical framework of Markov Decision Processes (MDPs). The defender strategy will be found using Reinforcement Learning. Lastly, the defender strategy will be learned and evaluated against other known strategies in a simulation environment.

Project L2: Designing and Simulating Cyber Attacks on Drones

Supervisors: Axel Andersson (axander@kth.se), György Dán (gyuri@kth.se)



The objective of this project is to design and implement various cyber attacks against autonomous vehicles (AV) such as drones or rovers. This could be done be leveraging open source simulation software such as the Gazebo Simulator. The aim of the project is to understand the vulnerabilities of autonomous vehicles and the impact these attacks can have.

Autonomous Vehicles rely on sensors, control algorithms and communication networks to navigate and perform other tasks. All of these components are vulnerable to cyber-attacks that can compromise the safety and security of these systems. The aim is to test these attacks on flight stacks that are used in real commercial drones/rovers such as PX4 or ArduPilot.

Main components of the Project

1. Set up simulation environment. Get comfortable with simulation software such as Gazebo and understanding key components of the PX4/ArduPilot flight stack.
2. Design a benchmark of missions. Develop a range of repeatable flight missions that can be analyzed later when different attacks are launched.
3. Attack Design. Develop different types of attacks such as:
    - False Data Injection (FDI). Different attacks could be constructed by injecting false data to certain sensors.
    - Communication Intercepting. Manipulate data that is exchanged between the AV and a central control system or data exchanged between AV's.
4. Attack Implementation. Implement the attacks by writing scripts and/or modifying the PX4/ArduPilot code and run simulations when the AV is being attacked.
5. Analysis. Analyze the outcome of the attack. What effect does the attacks have on the AV, do they cause the AV to fail its mission? Reason about potential counter-measures for the attacks.

If you are interested in cyber-security for control systems/autonomous vehicles and learning a new simulation tool then this project could be of interest to you! It is recommended to feel comfortable with programming, both PX4 and ArduPilot are large code bases written in C++.

Project L3: Adversarial Attacks on Multi-Agent Reinforcement Learning Systems

Supervisors: Kiarash Kazari (kkazari@kth.se), György Dán (gyuri@kth.se)

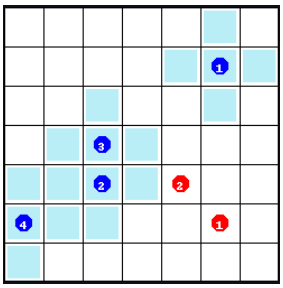


Figure 6: The Predator Prey OpenAI Gym Environment

In this project, the goal is to find effective attacks against multi-agent reinforcement learning (MARL) systems. A MARL system is a system where multiple autonomous agents operate in the same environment and want to accomplish some goal. If the agents work together as a team, it is called cooperative MARL (c-MARL). An example of this is a search-and-rescue task where multiple autonomous agents (e.g., drones) are searching for a target.

In the project we will study how robust MARL algorithms are by trying to construct attacks that cause the team of autonomous agents to fail their mission. A possible avenue could be to perturb the agents' observations, causing them to take bad decisions. One could consider both training and inference time attacks, that is, is the attacker present during the training of these algorithms or only when the agent is deployed in a mission.

Main Tasks of this Project

1. Choose an interesting and not too complex c-MARL task.
2. Design a simple and effective attack strategy against the MARL system that causes the system to fail the mission or decrease the general performance. This could be:
    - Perturbation Attacks: Introduce small changes in one or several of the agents observations, affecting decision making at inference-time or during training.
    - Behavioral Attacks: Introduce a malicious agent whose goal is to minimize the overall team reward.
3. Interact with open-source environments commonly used in RL and MARL research such as OpenAI's Gym or PettingZoo.
4. Train a MARL system to understand the baseline performance.
5. Implement the Attack Strategy.
6. Evaluation. Measure how the attacks affect the overall performance of the MARL system.
7. Report and Presentation: Document your findings and the process. Reason about potential countermeasures.

Tools we will use

Python and some automatic differentiation library like PyTorch, TensorFlow or Jax. OpenAI's Gym or PettingZoo for simulating an environment.

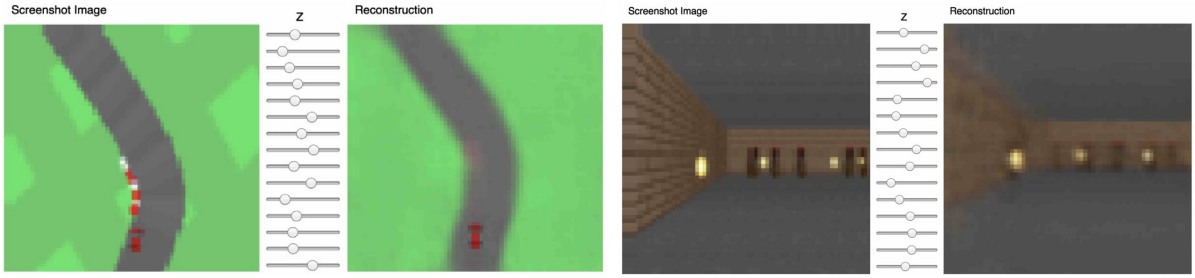
Related Work

Lin, Jieyu & Dzeparoska, Kristina & Zhang, Sai & Leon-Garcia, A. & Papernot, Nicolas. On the Robustness of Cooperative Multi-Agent Reinforcement Learning. SPW 2020, 62-68.

Kiarash Kazari, Ezzeldin Shereen, György Dán, ``Decentralized Anomaly Detection in Cooperative Multi-Agent Reinforcement Learning", in Proc. of Int. Joint Conf. on Artificial Intelligence (IJCAI), Aug. 2023

Project L4: Robustness of RNN-Based World Models

Supervisors: Manjodh Singh (manjodh@kth.se), György Dán (gyuri@kth.se)



World models are machine learning systems that learn to predict how an environment changes over time based on visual observations. They usually contain two parts: an encoder that compresses raw observations into a smaller latent state and a dynamics model that predicts how this latent state evolves when actions are taken. A controller is used to decide which actions to take by looking at the predictions of the model. Using such a learned model allows the agent to plan without interacting with the real environment at every step. However, this also raises concerns about robustness. Since the world model is an approximation of the real environment, the controller can be misled. An attacker could deliberately perturb the inputs to cause the agent to behave unsafely.

The goal of the project is to implement a small and interpretable world model, train it in a simple simulation environment, and then test how robust it is under attack. The main questions are how well the model can learn an accurate representation of the environment, how performance degrades under different kinds of perturbations, and how simple countermeasures can improve its robustness.

For more background information about the project, (Ha and Schmidhuber, 2018) and (Zeng et al. 2024) serve as good starting points for exploring the rich literature on world models.

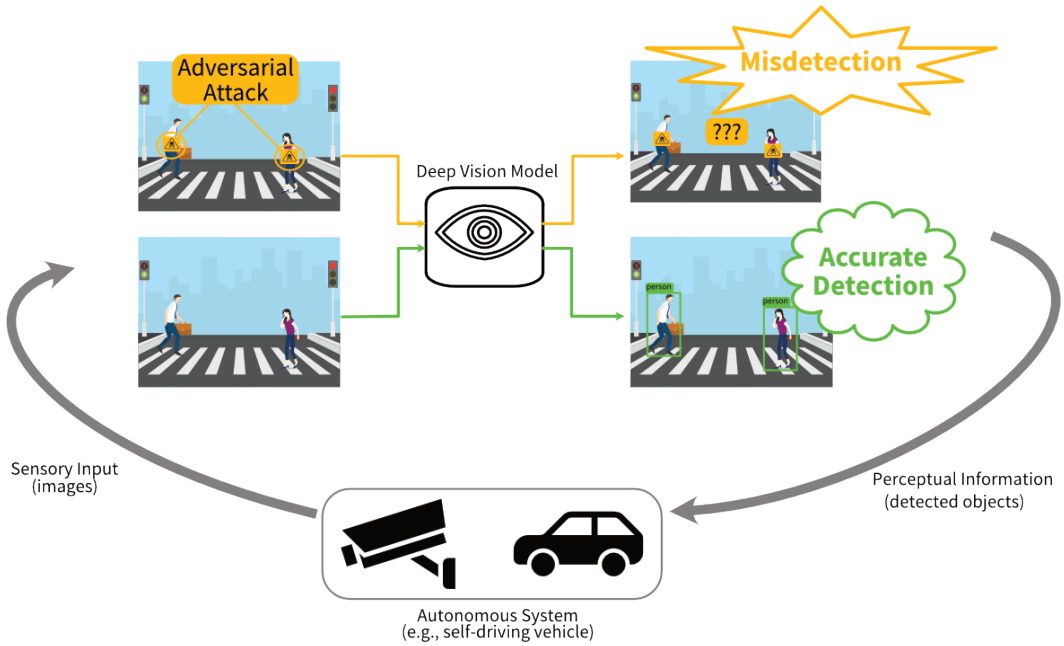
Related work:

David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In Advances in Neural Information Processing Systems 31, pages 2451–2463. Curran Associates, Inc., 2018. URL https://papers.nips.cc/paper/7512-recurrent-world-models-facilitate-policy-evolution. https://worldmodels.github.io.

Zifan Zeng, Chongzhe Zhang, Feng Liu, Joseph Sifakis, Qunli Zhang, Shiming Liu, and Peng Wang. World models: The safety perspective. In 2024 IEEE 35th International Symposium on Software Reliability Engineering Workshops (ISSREW), pages 369–376. IEEE, 2024.

Project L5: Adversarial Attack and Defense for Deep Vision Models

Supervisors: Mauricio Byrd Victorica (mbv@kth.se), György Dán (gyuri@kth.se)



Computer vision focuses on extracting high-level semantic information from visual data. Facial recognition for biometric authentication and dynamic obstacle detection in autonomous driving are examples of computer vision tasks.

Deep vision models are neural networks trained to perform such tasks, and they have been the dominant approach in the last two decades. In spite of their success, there are critical concerns regarding their deployment, especially in safety-critical scenarios. This project focuses on one critical vulnerability of deep vision models that has attracted a lot of interest in the last decade.

Visual adversarial attacks are perturbations designed to mislead a victim vision model into producing a specific output, for example, associating a person's face with the wrong identity [1], or failing to detect obstacles in the path of an autonomous car [2].

A lot of research has been devoted to developing defenses against adversarial attacks [3,4], but there are still many open questions that are relevant to crucial real-world computer vision scenarios. The objective of this project is to formulate and tackle one or more research questions that relate to the open challenges of adversarial attacks and defenses.

- Creating adversarial perturbations, i.e., attacks.
- Designing or improving mechanisms to defend ML models from adversarial attacks.
- Adapt existing attacks and defenses to unexplored scenarios (e.g., new victim models, attack formulations, datasets, etc.)

Related work:

[1] Mahmood Sharif et al. "Accessorize to a Crime: Real and Stealthy Attacks on State-of-the-Art Face Recognition", In Proc. of ACM CCS, 2016.
[2] Yi Zhu et al. "Malicious Attacks against Multi-Sensor Fusion in Autonomous Driving", in Proc. of ACM MobiCom, 2024.
[3] Mauricio Byrd Victorica et al., "SpaNN: Detecting Multiple Adversarial Patches on CNNs by Spanning Saliency Thresholds" in Proc. of SaTML, 2025.
[4] Mauricio Byrd Victorica et al., "Saliuitl: Ensemble Salience Guided Recovery of Adversarial Patches against CNNs", in Proc. of CVPR, 2025.

Project L6: Adding Static Information-Flow Checks to Secure AI Agents

Supervisor: Siyao Liu (siyaol@kth.se), Hamed Nemati (hnnemati@kth.se)

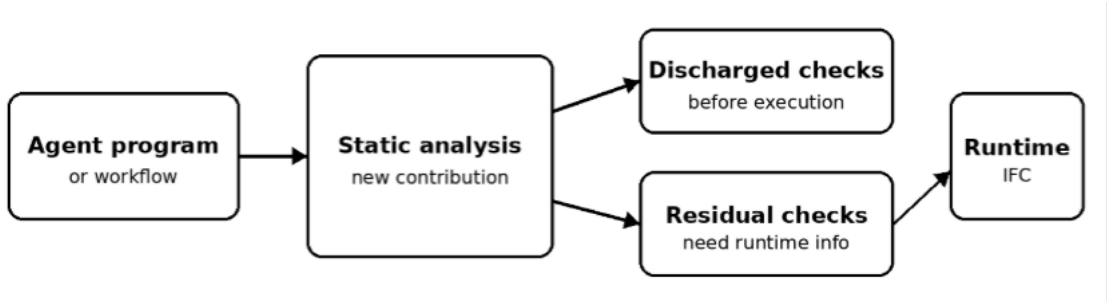


***Figure**: add static checks before execution while keeping uncertain cases protected at runtime.*

Large language models are increasingly used as AI agents that can read emails, browse the web, access files, and call external tools. This makes them useful, but also creates security risks: private information may be leaked, or untrusted content may influence an agent to perform an unsafe action.

Information-Flow Control (IFC) is a security technique for tracking how sensitive or untrusted information moves through a program. Recent agent-security systems such as CaMeL, FIDES, and LLMbda mainly enforce such policies while the agent is running [1-3]. In this project, you will extend one such system with a small static checker so that some checks can be decided before execution.

The project will extend an existing agent-security system with a simple conservative static information-flow analysis. Together with the supervisor, you will select a manageable class of information flows that can be checked before execution. Flows that cannot be decided safely in advance will continue to be protected by the existing runtime mechanism.

Main Tasks of this Project:

1. Learn the basics of IFC and security for tool-using AI agents.
2. Set up the selected existing system, such as CaMeL, FIDES, or LLMbda, and reproduce a few examples.
3. Implement static analysis for the selected class of information flows and integrate it with the existing system. The initial scope will be defined together with the supervisor and can be extended if time permits.
4. Combine static and runtime checking. Checks resolved by the analysis are handled before execution; uncertain cases remain protected dynamically.
5. Evaluate the result on representative agent programs or an existing benchmark such as AgentDojo [4]. Measure how many runtime checks can be resolved statically and explain the remaining limitations.

Tools and Background:

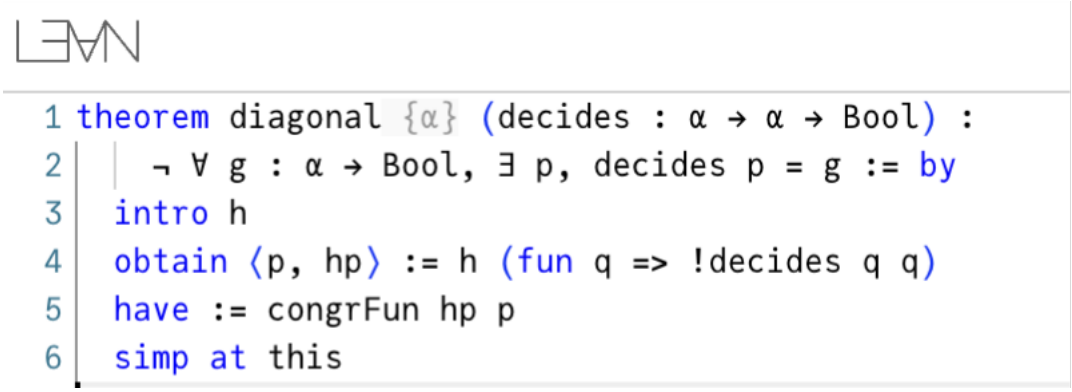
Good programming skills are recommended, but no previous knowledge of IFC, compilers, static analysis, or formal methods is required. The project has a clear core task, while the implementation and exact scope will be adapted to the students' background. CaMeL or FIDES offer a more implementation-oriented route; students interested in formal methods may instead use LLMbda and optionally explore Lean.

Related Work:

[1] Debenedetti, Edoardo, et al. "Defeating prompt injections by design." arXiv:2503.18813 (2025).
[2] Costa, Manuel, et al. "Securing AI agents with information-flow control." arXiv:2505.23643 (2025).
[3] Garby, Zac, Andrew D. Gordon, and David Sands. "The LLMbda calculus: AI agents, conversations, and information flow." arXiv:2602.20064 (2026).
[4] Debenedetti, Edoardo, et al. "AgentDojo: A dynamic environment to evaluate prompt injection attacks and defenses for LLM agents." NeurIPS 37 (2024): 82895-82920.

Project L7: Formal Verification in Cybersecurity

Supervisors: Yufei Liu (yufeiliu@kth.se), Hamed Nemati (hnnemati@kth.se)



```
theorem diagonal {α} (decides : α → α → Bool) :
    ¬ ∀ g : α → Bool, ∃ p, decides p = g := by
  intro h
  obtain ⟨p, hp⟩ := h (fun q => !decides q q)
  have := congrFun hp p
  simp at this
```

Figure: A machine-checked proof in Lean that illustrates the diagonal argument for the Halting Problem.

Formal verification aims to mathematically prove that a system (e.g. a piece of software) behaves as intended. Here, we model the program as a mathematical object, and its desirable properties are formulated as mathematical theorems. The process of formal verification becomes proving the theorems about these objects. Typically, this is done using programs called proof assistants, which automatically check the correctness of the proofs and can provide various degrees of automation in proof writing. Nowadays, with AI tools making code faster to write and bugs easier to find, this kind of guarantee is becoming more valuable.
In this project, you will understand both the strengths and limits of formal verification, and learn how to apply it in cases where correctness matters.

Possible directions include:

- Security protocols: model a real-world protocol (e.g. a VPN handshake or an IoT pairing process) and prove that it keeps secrets safe, or find an attack.
- Semantics modeling: give a precise definition of a protocol description language in a proof assistant (e.g. Lean or Isabelle/HOL) and prove properties about it.
- AI and verification: explore how to use AI tools to write machine-checkable proofs, or how to use verification to catch mistakes in AI-generated code.
- A topic of your choice within the scope of this project.

Useful background:
Interest in logic, security, or programming languages. **Curiosity** matters more than prior knowledge!

Project L8: Side Channel Analysis on Frodo Key Encapsulation Mechanism

Supervisor: Elena Dubrova, dubrova@kth.se, Electronics & Embedded Systems

With the probability of a party with access to quantum computing rising every year, the risk of using standard asymmetric encryption grows larger. Therefore, the use of post-quantum cryptographic algorithms is increasing to prevent any parties from decrypting sensitive information. The new algorithms open up for the need to test their susceptibility to side-channel analysis, so that awareness of their vulnerabilities can be spread.

The goal of the project is to find leakage in a masked implementation of FrodoKEM, using tools such as TVLA and deep learning to determine the susceptibility of the implementation to side-channel attacks.

In this project, you will:

1. Analyze the software implementation of the FrodoKEM algorithm to find potential leakage points
2. Analyze the leakage using statistical analysis.
3. Feed the data into a deep learning model to get a final idea of how susceptible the implementation is to side-channel attacks.



Context M: Information Engineering: Big Data & AI

Big data and artificial intelligence are broad topics with huge technological and economical potentials and therefore is of interest in many areas. From an engineering point of view, it is mostly related on how to process data. Currently, information and communication technology is penetrating all systems to make them *smart,* e.g. we envision smart cities, smart homes, smart grids, etc. or Internet of things in general. The smartness of the systems is built on the principle to sense the system environment and then draw smart decisions on it. However without algorithms that extract information from the data, the information is buried in the data and cannot be exploited. Thus, the process of extracting information will be the key ingredient of many future technologies and is the main objective of technologies nowadays known as artificial intelligence (AI), machine learning, data mining, pattern recognition, data analytics, adaptive signal processing etc. which are all instances of information engineering.

In general, we can say that the more data we have, the smarter the system will be. Thus, advanced smart systems sooner or later face the big data problem, which commonly means that the amount of the data is *too big* to be processed e.g. with standard tools. Therefore, there are huge research efforts developing novel information processing and data analytic methods, which enable future systems to deal with larger and larger data sets. More recently, agentic AI, i.e., AI systems that autonomously plan, call tools and interact with other agents, has emerged as a new paradigm that both creates new kinds of network traffic and helps automate the analysis of complex systems.

Innovative information processing and data analytic methods are traditional topics of the Information Science and Engineering Division. Thus, the sub-projects offered in this course will address fundamental topics and problems in the area with a strong engagement of the division's teachers. Accordingly, all offered Bachelor projects are closely related to on-going research projects in the division. This year, we offer projects on agentic AI in 5G/6G networks, an open WiFi stack for software-defined radios, restoration of scanned historical documents, deep learning-based biometric identification, data-driven state estimation for 6G sensing, and signalling strategies in decentralized control. Since information processing is quite abstract, all projects require a good mathematical background and solid programming skills.

Project M1: Agentic AI as a New Traffic Type in 5G/6G Networks

Contact: Wilson Nelson, wilsonan@kth.se, Division of Information Science and Engineering

Background

AI agents behave differently from conventional applications such as video streaming, web browsing, and chatbots. They make many API and tool calls, fetch web pages, communicate with other agents through dedicated protocols, and upload content such as images and documents. As a result, their traffic is expected to differ significantly from that of these applications.

In this thesis, real AI agents are run over a 5G protocol stack, initially using OpenAirInterface's RF simulator (OAI rfsim) and the agentic rfsim framework developed by our research group. Because the complete stack is visible, agent behavior can be linked to both IP-level and MAC-level metrics. Analyzing and modeling this traffic will help prepare for the 6G era, in which agentic traffic is expected to become ubiquitous.

Problem Statement and Research Questions

There is currently little research on AI agent traffic and how it interacts with the 5G protocol stack. It is therefore unclear whether today's 5G scheduling, configuration, and QoS mechanisms serve agents well, or whether the network can even recognize agent traffic in order to treat it differently. This thesis addresses the following research questions:

1. How does the traffic of autonomous AI agents differ from that of other applications, at both IP level and 5G MAC level?
2. How can agent traffic be modeled so that the model reproduces its packet-level statistics under controlled conditions?
3. Can agent traffic be detected and distinguished from other application traffic, such as video streaming and web browsing?

Proposed Approach

1. **Literature review:** Study AI agent concepts and protocols, 5G uplink scheduling, and traffic modeling.
2. **Experimental setup:** Build an AI agent whose traffic is routed through the rfsim 5G connection. The setup must capture packet traces, agent event logs, and gNB and UE logs with synchronized timestamps.
3. **Data collection:** Design and run experiments to collect a dataset of agent traffic, together with traffic from baseline applications for comparison.
4. **Traffic characterization:** Characterize and analyze agent traffic using the collected IP-level and MAC-level traces.
5. **Traffic modeling:** Based on the analysis, develop a traffic model that explicitly captures the key characteristics of agent traffic.
6. **Agent traffic detection:** Use the collected dataset to develop a classifier that distinguishes agent traffic from other application traffic. Analyze which features drive the detection and whether they generalize to other agentic traffic scenarios.
7. **Thesis writing and presentation:** Write the thesis and prepare the presentation based on the results of the previous steps.

Required Skills

- Good knowledge of **statistics** and programming in **Python**.
- Experience with **software development** and **5G concepts**.
- Willingness to learn new concepts, in particular agentic AI and traffic modelling.

References

[1] Ericsson (2026). **Ericsson Mobility Report, June 2026: AI will drive future uplink**.

[2] Kundu, L., et al. (2026). **AI-Native 6G for Distributed Intelligence: Traffic Characteristics, Awareness, and AI Grid**. *arXiv preprint* arXiv:2608.14627.

[3] Montieri, A., Nascita, A., and Pescapè, A. (2026). **From prompts to packets: A view from the network on ChatGPT, Copilot, and Gemini**. *Computer Networks*, 112237.

Project M2: Autonomous Interventional Causal Discovery of 5G Latency

Contact: Wilson Nelson, wilsonan@kth.se, Division of Information Science and Engineering

Background

End-to-end latency in 5G depends on many interacting factors, such as numerology, TDD pattern, scheduling, HARQ retransmissions, and buffering. In practice, it is difficult to determine quantitatively which of these parameters actually drive latency in a given situation. A causal graph that captures how these variables influence each other and latency would go a long way toward resolving this.

Causal graphs are usually learned with causal discovery algorithms that work on observational data. These algorithms rely on statistical dependencies in the data, which cannot always distinguish true cause-and-effect relationships from correlations. Interventional experiments, where one variable is deliberately changed while others are held fixed, give much stronger evidence of causation. However, designing interventional scenarios and running them across many potentially dependent parameters takes significant time, compute, and effort.

Recent advances in agentic AI make it possible to automate such workflows. An agentic framework that autonomously designs and runs interventional experiments, and derives a causal graph from the results, would greatly reduce this effort. This thesis is a first step toward that vision, using the 5G latency of an OpenAirInterface rfsim setup as the starting point.

Such a system would have several benefits. It can validate and strengthen causal graphs obtained from observational data. It can provide a knowledge base for root cause analysis of anomalies observed in 5G systems. It can also support the development and validation of 5G digital twins, and help quantify the sim-to-real gap when moving from simulation to real deployments.

Problem Statement and Research Questions

Causal discovery from observational data alone is often unreliable, while manual interventional studies are slow and hard to repeat on a different simulator or testbed. To address this gap, this thesis develops a tool that autonomously discovers the causal structure of 5G latency by designing and running its own interventional experiments. The tool builds on agentic AI, which can automate the experimental workflow and focus on an informative subset of the large space of possible experiments instead of testing every combination.

The thesis addresses the following research questions:

1. How can an agent efficiently search the experiment design space to select meaningful interventional experiments, and how many experiments does it need to arrive at a reliable causal graph?
2. How does the causal graph discovered by the agentic framework compare with graphs obtained from standard causal discovery algorithms that rely on observational data?

Proposed Approach

1. **Literature review:** Study relevant material on causality, causal discovery algorithms, agentic AI, and agentic frameworks.
2. **Scope definition:** Define the variables to consider, the architecture of the framework, and the evaluation criteria.
3. **Familiarization:** Get familiar with the existing agentic framework (developed by our group) for the OpenAirInterface rfsim simulator.
4. **Design and implementation:** Design and build the autonomous interventional framework on top of the existing rfsim agentic framework.
5. **Validation:** Test and refine the framework on known interventional scenarios with expected outcomes.
6. **Causal discovery:** Run autonomous interventional experiments to obtain causal graphs for latency.

7. **Benchmarking:** Compare the obtained causal graphs with those from standard causal discovery algorithms and the 3GPP knowledge base.
8. **Thesis writing and presentation.**

Required Skills

- Good knowledge of **statistics** and programming in **Python**.
- Experience with software development, 5G concepts, or agentic AI is a plus.
- Willingness to learn new concepts, in particular causality and causal inference.

References

[1] Havrilla, A., Alvarez-Melis, D., and Fusi, N. (2025). **IGDA: Interactive graph discovery through large language model agents**. *arXiv preprint* arXiv:2502.17189.
[2] **LLM-Powered Agentic AI for 5G/6G Networks: A Tutorial and Survey** (2026). *arXiv preprint* arXiv:2607.16066.
[3] Zhang, K., et al. (2021). **gCastle: A Python Toolbox for Causal Discovery**. *arXiv preprint* arXiv:2111.15155.

Project M3: A WiFi Network Stack from rfSim to the ExPECA Testbed

Contact: Wilson Nelson, wilsonan@kth.se, Division of Information Science and Engineering

Background

WiFi (IEEE 802.11) is among the most widely deployed wireless access technologies, connecting a large share of today's smartphones and laptops to the Internet. However, its protocol implementation typically resides in proprietary chipsets, which limits the extent to which it can be inspected or modified for research purposes. Software-defined radios (SDRs) address this limitation by implementing radio functionality in software running on general-purpose processors. A software implementation of the WiFi protocol stack thus exposes every processing stage, from the bits of a packet to the transmitted waveform, to observation, measurement and experimentation.

In cellular research, such software-based development is already well established. OpenAirInterface (OAI) provides an open-source implementation of 4G and 5G protocol stacks in which the same code can first be validated with the RF simulator (rfSim) on a single computer and then deployed unchanged on SDR hardware.

In this thesis, the aim is to design and implement a complete WiFi protocol stack and to evaluate it first in simulation and then over the air in the ExPECA testbed that uses SDR hardware. The work combines software development with signal processing concepts and experimental SDR research, which are central to the development of future wireless systems such as 6G. The resulting stack will serve as an open platform for education and further research, for instance on novel medium access control (MAC) schemes and AI-assisted protocol development. In addition, comparing simulation with over-the-air measurements will quantify how well simulation predicts the performance of a real wireless link.

Problem Statement, Aim and Research Questions

Problem. OAI's rfSim and SDR testbeds such as ExPECA are built around 4G/5G. There is no WiFi stack that uses OAI's radio interface, so WiFi experiments cannot use the same fast simulation-to-testbed workflow. It is also unclear how well a WiFi PHY/MAC verified in rfSim performs on real SDRs, where carrier frequency offset, hardware impairments and real-time limits come into play.

Aim. Develop and evaluate a WiFi (802.11a-based) network stack that runs unchanged in OAI rfSim and on ExPECA's USRP E320 radios.

The thesis addresses the following research questions:

1. How can a WiFi (802.11a-based) stack be built so that the same implementation runs in OAI rfSim and on USRPs?
2. How well does the stack perform in simulation?
3. How well does it perform over the air in ExPECA, and what explains the difference from simulation?

Proposed Approach

1. **Literature review and learning foundations:** Learn the IEEE 802.11a/g PHY and MAC, existing SDR WiFi implementations, OAI's architecture and rfSim, and the ExPECA testbed.
2. **Familiarization with the OAI codebase:** Build OAI and run its 5G base station and UE together in rfSim to understand the simulation workflow.
3. **Investigation of the radio layer:** Investigate how to send and receive samples through OAI's radio interface and rfSim, and build a simple test link between two nodes.
4. **Develop the PHY layer:** Implement the 802.11a transmitter and receiver step by step, checked against a reference model and the standard's worked example.
5. **Develop the MAC layer:** Add 802.11 frames, error checking, acknowledgements and retransmission.

6. **Evaluate in simulation:** Measure error rates and throughput in rfSim under different channel conditions.
7. **Move to ExPECA:** Learn the testbed's reservation and container workflow, find out what is needed to drive its USRP radios, and make the first over-the-air transmissions.
8. **Evaluate over the air:** Measure the link on different radio pairs and compare with the simulation results.
9. **Thesis writing and presentation:** Write the thesis, share the code and experiment scripts as open source, and prepare the presentation.

Most of the above steps and software modules should be accompanied by appropriate unit tests to facilitate debugging. The software stack is intended to be developed in a modular manner so as to understand and learn component-level behavior.

Required Skills

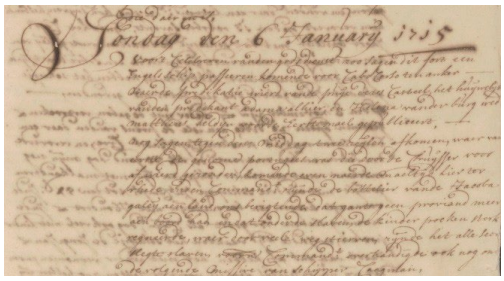
- Good programming skills in **C** (or C++), and **Python** for scripts and data analysis.
- Comfortable with **Linux**, the command line and git, and willing to work in a large open-source code base.
- Basic knowledge of **digital communications and signal processing** (modulation, OFDM, FFT) and of **computer networks** (protocol layers, MAC).
- Willingness to learn about wireless systems and patience for systematic debugging.
- Nice to have: experience with SDRs and Docker/containers.

References

[1] **IEEE Std 802.11-2020, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications**, clause 17 (OFDM PHY) and the annex with the worked frame-encoding example.
[2] Mostafavi, S., et al. (2023). **ExPECA: An Experimental Platform for Trustworthy Edge Computing Applications**. *ACM/IEEE Symposium on Edge Computing (SEC)*.
[3] **OpenAirInterface 5G source code**.
[4] **OAI RF simulator (rfSim) documentation**.

Project M4: Improving Scanned Documents where the Reverse Side Shows Through

Contact: Mats Bengtsson, ***matben@kth.se***, Division of Information Science and Engineering



Background

Even though the internet is flooded by more or less interesting data, generated from all sorts of sources, ranging from artists and bloggers to various sensors and other machines, a lot of valuable information is still buried in libraries and archives around the world. This cultural heritage, in the form of printed books and hand written manuscripts is successively made available in digital form, but this digitalization involves many challenges.

One of the challenges when scanning a document appears if the paper is so thin that the text or images from the reverse side of the page shows through. Especially, in hand-written manuscripts and some old prints, the ink may have bled through the paper, making it more or less impossible to interpret the scanned information. However, show-through effects can also appear when scanning modern printed papers, since some light passes through the paper during the scanning process.Luckily, you often have access to scans of both sides of the paper, meaning that you have access to two images, which both contain information about the original information that was written or printed on the two pages. This makes it possible to apply different signal processing methods to try to separate the information from the two pages. The bleed-through or show-through phenomenon is typically non-linear, thus requiring non-linear restoration methods. Related techniques for separating multiple sources that have undergone a mixing process are also used in audio processing, for separation of several simultaneous sound sources, based on microphone signals. This is called blind source separation. An additional complication is that of aligning the scans of the two sides of the page.

Aim of the Thesis Project

The goal of this project is to identify a couple of algorithms from the literature and implement and evaluate these both on data that has been generated synthetically, so that the “true” solution is known, but also on authentic scans. Even though neural networks are known to be good candidates for non-linear processing algorithms, there are surprisingly few publications proposing such solutions for this particular problem. Therefore, there might be possibilities to propose and try out new solutions. The evaluation of the different algorithms can be based both on some mathematically defined “objective” performance criterion, but also on a more “subjective” visual impression.

References

[1] Savino, P., and Tonazzini, A. (2023). **Mathematical models and neural networks for the description and the correction of typical distortions of historical manuscripts**. *International Conference on Computational Science and Its Applications*. Springer Nature Switzerland.

[2] Oda, M., and Miyajima, H. (2014). **Show-through cancellation in scanned documents using two-layer bidirectional neural networks**. *Joint 7th Int. Conf. on Soft Computing and Intelligent Systems (SCIS) and 15th Int. Symp. on Advanced Intelligent Systems (ISIS)*. IEEE.

[3] Tonazzini, A., Salerno, E., and Bedini, L. (2007). **Fast correction of bleed-through distortion in grayscale documents by a blind source separation technique**. *Document Analysis and Recognition*, vol. 10, pp. 17–25.

[4] Sharma, G. (2001). **Show-through cancellation in scans of duplex printed documents**. *IEEE Transactions on Image Processing*, 10(5), 736–754.

[5] Ehsandoust, B., et al. (2017). **Blind source separation in nonlinear mixtures: separability and a basic algorithm**. *IEEE Transactions on Signal Processing*, 65(16), 4339–4352.

Project M5: Data-driven State Estimation of Model-free Dynamical Processes

Contact: Saikat Chatterjee, *sach@kth.se*, and Fredrik Cumlin, *fcumlin@kth.se*, Division of Information Science and Engineering

Background

The Apollo program, which successfully landed the first humans on the moon, is still considered one of mankind's greatest achievements. Among the technological innovations that played a role in the success of the Apollo program is a filtering method based on the extended version of an algorithm developed by Rudolf Kalman in the late 1950s, which was used to track and estimate the trajectory of the spaceship. This algorithm, known as the Kalman Filter (KF), and its extension into the extended KF (EKF), provide an accurate and reliable real-time trajectory estimation while still being simple to implement and applicable on the hardware-limited navigation computer of the Apollo spaceship. To date, the KF is among the most celebrated algorithms in signal processing and electrical engineering at large, with numerous applications, including radar, biomedical systems, vehicular technology, navigation, and wireless communications [1].

KF, its extensions, and particle filters (PFs) are part of the larger topic called Bayesian state estimation (BSE). BSE of an underlying stochastic and complex dynamical process from (noisy) linear and nonlinear measurements is a fundamental problem in signal processing, machine learning and automatic control, underpinning applications such as tracking, navigation, and time-series analysis. For example, tracking a drone and a drone swarm using radar measurements. Traditional BSE methods, such as KF, its extensions and PFs are model-driven. All of these state estimation methods use a-priori knowledge of the underlying dynamical system. For example, tracking a drone using KF requires knowing a model of how the drone flies.



Problem Statement

The big issue is what happens if we do not have a clear knowledge of the drone's flight capability and dynamics. We may not know how many rotors the drone has, what its power, mass, velocity and acceleration capability, direction change capability, etc. are. Then, we do not have a proper model of the drone, and it is difficult to use the KF or any model-driven BSE method. Here comes the approach of developing data-driven state estimation methods for a model-free dynamical process. We try to learn everything from a large amount of (noisy) measurement data collected by the related sensors, and track the drone(s). The sensors are radars and base stations. This project will address the design of such data-driven methods. The project will be based on the recent method called Data-driven Nonlinear State Estimation (DANSE) developed in Saikat Chatterjee's research group at ISE, EECS [2].

The project is part of a larger collaboration between KTH and Ericsson for the development of an Integrated Sensing and Communication (ISAC) 6G system.

References

[1] Shlezinger, N., Revach, G., Ghosh, A., Chatterjee, S., Tang, S., Imbiriba, T., Dunik, J., Straka, O., Closas, P., and Eldar, Y. C. (2025). **Artificial Intelligence-Aided Kalman Filters: AI-Augmented Designs for Kalman-Type Algorithms**. *IEEE Signal Processing Magazine*.

[2] Ghosh, A., Honoré, A., and Chatterjee, S. (2024). **DANSE: Data-driven Non-linear State Estimation of Model-free Process in Unsupervised Learning Setup**. *IEEE Transactions on Signal Processing*, vol. 72, pp. 1824–1838.

Project M6: Low-Power Signalling Strategies for Witsenhausen's Counterexample

Contact: Mengyuan Zhao, *mzhao@kth.se*, and Tobias Oechtering, *oech@kth.se*, Division of Information Science and Engineering

Problem Statement

How should a controller act when its actions both influence a system and communicate information to another controller? This question is central to decentralized control, where decision-makers have access to different information. Witsenhausen's counterexample is a classical problem in distributed decision-making demonstrating that nonlinear strategies can outperform linear strategies, even in a system with Gaussian noise and quadratic costs [1].

The problem involves two decision-makers. The first observes an initial state and modifies it through a control action, which incurs a power cost. The second observes a noisy version of the modified state and attempts to estimate it, thereby incurring an estimation cost. The challenge is to balance the power used by the first decision-maker against the estimation error of the second.

The recent work Low-Power Optimal Strategy for Witsenhausen Counterexample [2] introduces a low-power estimation (LoPE) strategy. It uses a step function for the first control action, making the resulting system state a piecewise-linear function of the initial state with slope one within each segment. This approach contrasts with Witsenhausen's two-point strategy as well as its n-point generalization [3], which design the target system state to be a step function instead. This motivates investigating more general piecewise-linear strategies: can adjusting the slopes, thresholds, and offsets improve performance when only a small amount of control power is available?

Aim of the Thesis Project

The aim is to implement, optimize, and compare signalling strategies for Witsenhausen's counterexample, with particular emphasis on piecewise-linear strategies in the low-power regime. Specifically, the students will:

1. **Study the problem and existing strategies.** Understand the roles of control and implicit communication, and review linear strategies, classical quantization-based strategies, and the LoPE approach.
2. **Implement a numerical evaluation framework.** Compute control power and estimation error using the minimum mean-square error estimator for each strategy. Validate the implementation against analytical benchmarks and selected results from the reference paper.
3. **Explore piecewise-linear strategies.** Starting from simple symmetric mappings with a small number of segments, optimize their slopes, thresholds, and offsets. The mappings may include jumps between segments, as in LoPE and sawtooth-shaped strategies.
4. **Compare performance under common power budgets.** Investigate how estimation error changes with the available power, noise level, and number of segments. Compare against the best linear strategy and LoPE, and include classical quantization-based strategies where the power budget permits.
5. **Interpret the findings.** Examine whether adjustable slopes provide meaningful improvements over fixed-slope LoPE and how the optimized strategy changes as the power budget decreases. Assess numerical accuracy and sensitivity to optimization initialization.

The expected outcome is a reproducible simulation framework and a comparative analysis identifying when additional strategy complexity is beneficial. The core project is numerical and suitable for two students completing a 15-credit bachelor's thesis. An optional extension is an analytical study of a simple piecewise-linear strategy near zero power.

Required Skills

- Solid knowledge of **probability** and **calculus**.
- Programming experience in **Python** or **MATLAB**.
- Interest in optimization, estimation, and communication or control theory.

References

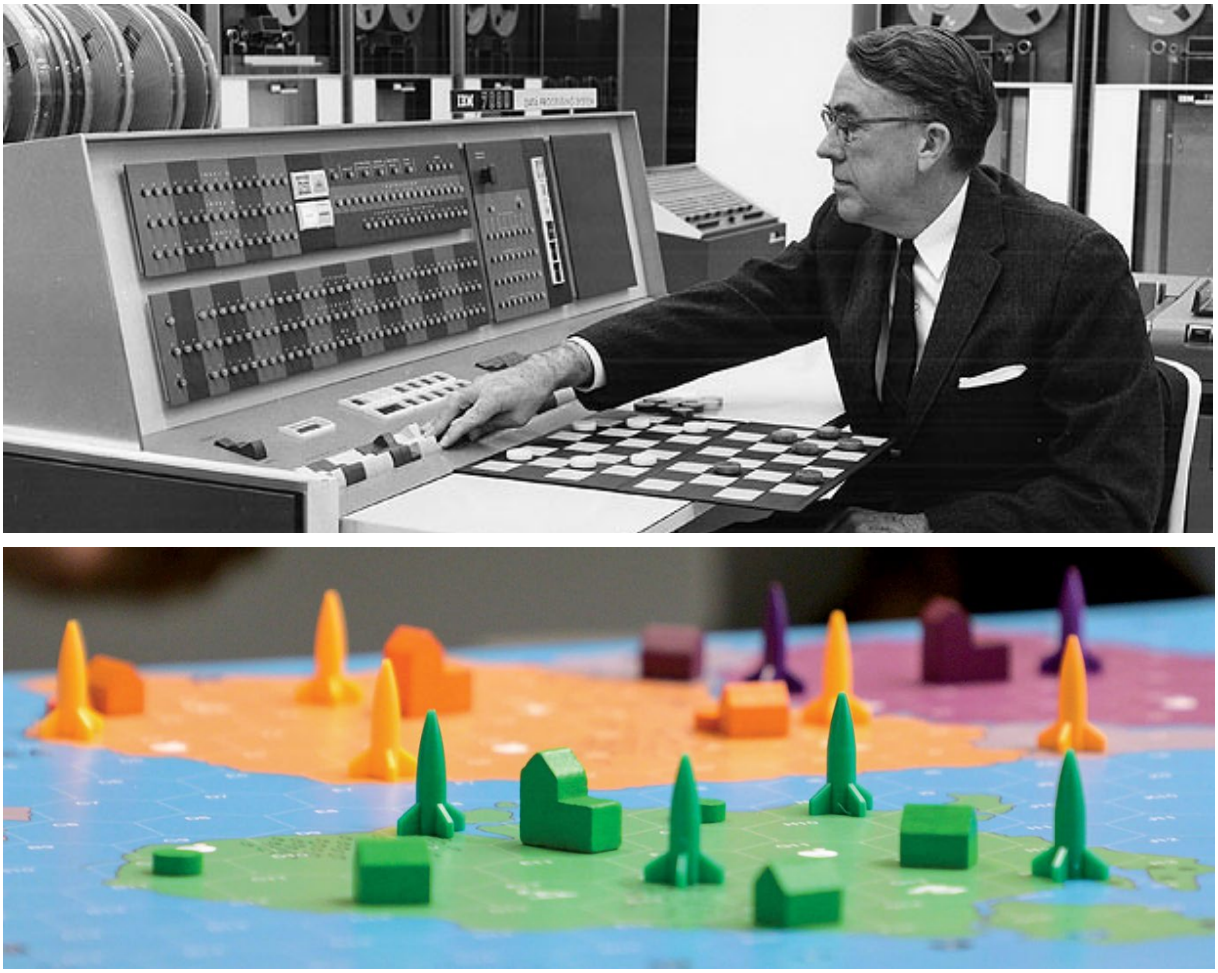
[1] Witsenhausen, H. S. (1968). **A Counterexample in Stochastic Optimum Control**. *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147.

[2] Zhao, M., Le Treust, M., and Oechtering, T. J. (2025). **Low-Power Optimal Strategy for Witsenhausen Counterexample**. *2025 Conference on Decision and Control (CDC)*, pp. 4503–4509.

[3] Lee, J., Lau, E., and Ho, Y.-C. (2001). **The Witsenhausen counterexample: a hierarchical search approach for nonconvex optimization problems**. *IEEE Transactions on Automatic Control*, vol. 46, no. 3, pp. 382–397.

Context N: AI, games, and strategy

Context Responsible: Mika Cohen (mikac@kth.se)
Division of Theoretical Computer Science



Top: The very first AI-program was a computer program that played Checkers (www.chessprogramming.org/Christopher_Strachey). Bottom: The strategy game SIGNAL is used to study nuclear escalation dynamics (pong.berkeley.edu/e-game/).

AI and strategy games have been intimately connected since the very inception of AI. Classic strategy games such as Checkers, Chess, Go, and Poker have served as shared research goals within the AI-community and as benchmarks for evaluating AI-algorithms.

The interest in strategy games within AI has also been motivated by the prospect of supporting real world decision making and strategy. In the social sciences and elsewhere, there is a long tradition of analyzing human interactions of various kinds as strategy games – from relatively peaceful interactions such as business negotiations or stock trading to directly hostile interactions such as dogfights in air combat.

Given a model of a conflict as a strategy game, the strategy space can be explored systematically. However, exploring the strategy space manually, e.g., by repeatedly playing the game, requires considerably time and effort. The recent advances in AI for strategy games have opened up for the possibility of automating decision making and strategy development. Application areas range from security and defence to finance

Project N1: Knowledge in Multi-Agent Games

Supervisor: Dilian Gurov (dilian@kth.se)
Division of Theoretical Computer Science



Key Words: Multi-Agent Systems, Game Theory, Knowledge-Based Strategies

In a multi-player game, a coalition of **players** (also called agents) is attempting to achieve an **objective** within a (potentially hostile) environment, considered to be the opponent. Solving such a game means to find a **strategy** that achieves the objective regardless of the moves of the environment. Rescue missions involving robots and humans or pursuit-evasion games are examples of such games, often called multi-agent systems.

An interesting, but complicating circumstance is when the players have limited information about the current state of affairs, say due to limited observation capabilities. Such games are called games of **imperfect information**. A related aspect is posed by the communication capabilities between players. The problem of strategy synthesis under imperfect information and limited communication is known to be hard, and is an active research area. The present project investigates the modelling of such games, as well as algorithmic and machine learning-based techniques for strategy synthesis. In particular, the project focuses on strategies based on the notion of **knowledge**. In the context of this project, knowledge refers to information, structured suitably, stored and updated during the course of a play, for deciding on a course of action. Especially interesting is **higher-order knowledge**, where players maintain and use during play knowledge about the other players' knowledge.

Inspirational Reading:

[1] Gurov, D., Goranko, V., Lundberg, E.: *Knowledge-Based Strategies for Multi-Agent Teams Playing Against Nature*. Artificial Intelligence, vol.309, 2022, DOI: 10.1016/j.artint.2022.103728

[2] Doyen, L., Raskin, J.F.: *Games with imperfect information: Theory and algorithms*. Lectures in Game Theory for Computer Scientists pp. 185–212 (2011)

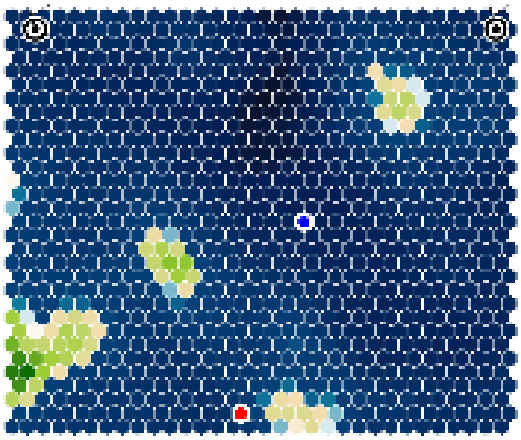
[3] Berwanger, D., Kaiser, L., Puchala, B.: A *perfect-information construction for coordination in games*. In: Foundations of Software Technology and Theoretical Computer Science (FSTTCS'11). LIPIcs, vol. 13, pp. 387–398 (2011)

[4] Huang, X., van der Meyden, R.: *Synthesizing strategies for epistemic goals by epistemic model checking: An application to pursuit evasion games*. In: Proceedings of AAAI 2012 (2012)

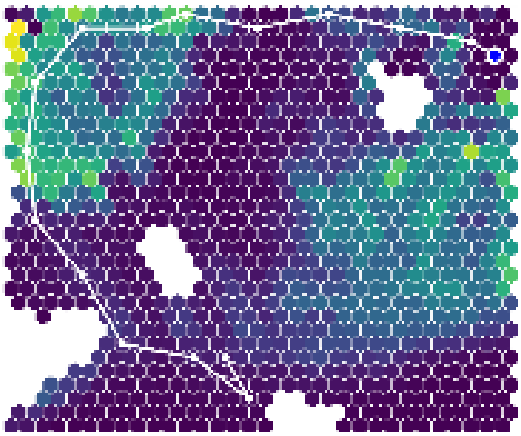
Project N2: Anti-Submarine Warfare with Counterfactual Regret Minimization

Supervisor: Mika Cohen (mikac@kth.se), Division of Theoretical Computer Science, Christopher Limér (christoffer.limer@foi.se), FOI

Summary: In this project, you will apply a poker-playing algorithm to the problem of planning Anti-Submarine Warfare in an optimally unpredictable way.



How should blue pursuit red?



Randomized search route

Background

A surface vessel (marked in red on the map to the left) pauses in midwater, an indication it might be launching an unmanned underwater vehicle. There are two critical infrastructure assets (marked as flags) that need to be protected from underwater threats. A patrol vessel with a dipping sonar is dispatched to the area, arriving after 1 hour to the position (marked in blue). How should the patrol vessel search?

In the Anti-Submarine Warfare (ASW) planning problem above, the patrol vessel needs to be unpredictable, or else the adversary will soon learn how the vessel tends to patrol (by studying the patrolling over time by means of satellite, passive sonar, radar, etc). Indeed, provoking patrols in order to study them is a standard scouting procedure also in naval grey zone conflicts.

It has long been assumed that unpredictability in ASW requires psychological intuition and cunning. Today, however, computer algorithms can outperform human players in challenging bluffing games. Most notably, bots are now unbeatable in poker, including the most demanding forms of poker. This invites a compelling question: Could poker-playing algorithms also play a game of hide-and-seek with a submarine?

Purpose

In this project, you will apply a poker-playing algorithm to the problem of planning ASW in an unpredictable way. More specifically, you will apply *Counterfactual Regret Minimization* (CFR) self-play, the basic building block of super-human poker-bots, to the *Princess and Monster game*, an abstraction of ASW studied in operational analysis.

Method

The CFR self-play algorithm is adapted to a variant of the Princess and Monster game that cannot be solved analytically. As a first step, a standard version of CFR self-play is implemented. Subsequently, CFR self-play is optimized to the specifics of the Princess and Monster game using e.g. Fixed-Strategy Iteration CFR or some other technique suggested by the project participants. The Princess and Monster game is implemented in python, but the programming language for implementing CFR can be chosen freely.

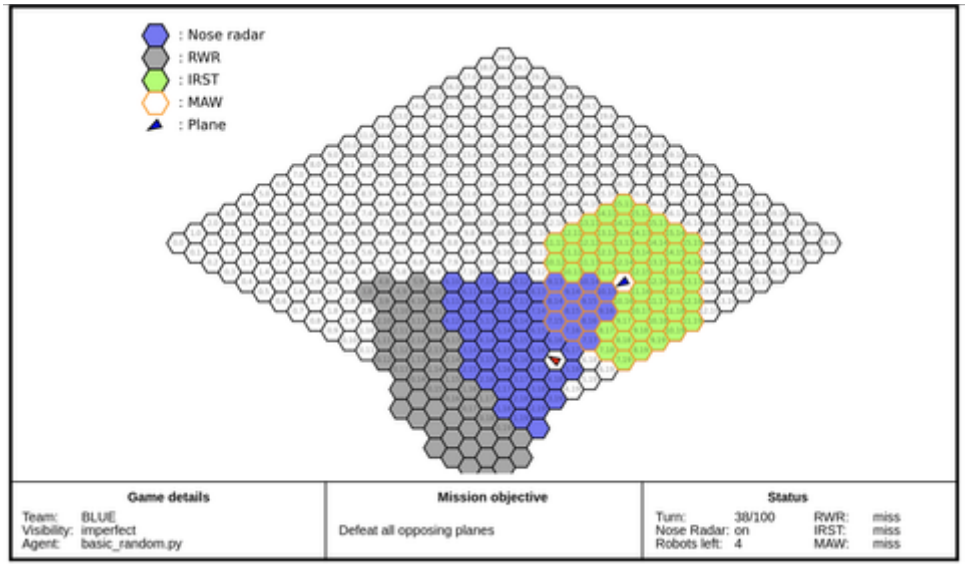
Context
Self-play techniques from poker are being applied to tactical decision support for ASW at the Swedish Defence Research Agency (FOI). The map on the righthand side above shows a solution the above ASW planning problem generated with a prototype tactical decision support. The search route (in white) is sampled from an optimally randomized search tactics, computed using game-theoretically sound self-play techniques from poker-bots. The heatmap shows the probability distribution of the submarine's location based on the successive sonar reports received so far (brighter green/yellow indicating higher probability). For more details, see the first reference below. The self-play techniques used there are rather complex, in this project, however, you will apply only CFR, a simpler, core component of poker-bots.

Further reading
Anti-Submarine Warfare Planning Using Public Belief States and Self-Play, Christoffer Limér et al., International Conference of Machine Learning and Applications (ICMLA), 2026.
Steps to building a Poker AI, Thomas Trenner, https://medium.com/ai-in-plain-english/steps-to-building-a-poker-ai-part-1-outline-and-history-58fbedaf6ded

Project N3: Beyond Visual Range Air Combat with Monte-Carlo Tree Search

Supervisor: Mika Cohen (mikac@kth.se), Division of Theoretical Computer Science, Felix Rydell (felix.rydell@foi.se), FOI

Summary: In this project, you will apply Monte-Carlo Tree Search to generate tactics for beyond visual range air combat.



How can blue outmanoeuvre red?

Background. In early 1941, when German submarines were destroying Allied shipping at a devastating rate, Churchill ordered the British Navy to "Find out what's going on and sink the U-boats". A new tactics development unit was created, the Western Approaches Tactical Unit (WATU), where staff simulated submarine attacks and developed countermeasures through wargaming. The rules of the war games reflected known physical properties of merchant ships, escorts and submarines in terms of speed, turning circle, visibility, armament and so on, but the rules left tactical decisions about formation, etc. open to the players to choose freely. Experimenting with tactics, the staff arrived at the best formations and search patterns for protecting the convoys through a creative, iterative process of trial-and-error.

The explorative wargaming at WATU played an important role for the development of the Battle of the Atlantic. Today it is not unreasonable to assume that “Similar challenges in the future could be tackled even more quickly and effectively with the help of AI programs like AlphaZero”[1], thereby “blurring the boundary between [professional] wargaming, game theory and [military] operational analysis”[2].

AlphaZero. AlphaZero is a general AI for double-sided strategy games that learns to play a given game through massive amounts of self-play. As AlphaZero repeatedly plays the game against itself, its decision-making converges towards the game-theoretic optimal play. Published already back in 2018, AlphaZero is still considered one of DeepMind's flagships. AlphaZero has helped to change the understanding of classic strategy games such as chess and go, uncovering more effective tactics that have eluded centuries of human creativity. More recently, AlphaZero, in combination with large language models, is being used for code generation and as such has discovered novel algorithms for ubiquitous computational tasks that surpasses the existing state-of-the-art.

Micro-wargames with AlphaZero. The anti-submarine warfare games in WATU belong to a type of small-scale wargames, sometimes called micro-wargames, which are in many ways reminiscent of classic strategy games such as Go and Chess. The difference is that the game rules aim to correctly reflect the dynamics of a certain real military conflict situation, albeit at a high level of abstraction.

Since AlphaZero has been shown to master classic strategy games and provide a window towards optimal gameplay in them, it might be expected that AlphaZero will similarly be able to provide a window towards optimal gameplay in micro-wargames as well.

Purpose
The project aims to apply Monte-Carlo Tree Search (MCTS), a core building block of AlphaZero, to generate tactics for a simple micro-wargame that models modern beyond visual range air combat.

Method
The MCTS self-play algorithm is adapted to the BVR micro-wargame (written in python, screen shot above). As a first step, a standard version of MCTS self-play is implemented. Subsequently, MCTS self-play is optimized to the fog-of-war aspects of the BVR game using, e.g., Information Set Monte Carlo Tree Search.

Context
AlphaZero and related self-play techniques are being applied to micro-wargames at the Swedish Defence Research Agency (FOI). For more details, see the first reference below.

Further reading
Explorative Wargaming with a Superhuman Tactician in the Team: A Controlled Experiment, J. Brynielsson et al., International Command and Control Research and Technology Symposium, 2026
The Animated Monte-Carlo Tree Search (MCTS), Kurbiel, https://medium.com/data-science/the-animated-monte-carlo-tree-search-mcts-c05bb48b018c.

Project N4: Explainable Fuzzy Challenge 2027

Supervisor: Mika Cohen (mikac@kth.se), Division of Theoretical Computer Science, Christopher Dristig Stenström (christopher.stenstrom@foi.se), FOI

Sammanfattning: I det här projektet spelar du det klassiska arkadspelet Asteroids med förklarbar AI (XAI).



”I det klassiska arkadspelet Asteroids manövrerar tvådimensionella rymdskepp för att undvika kollisioner med asteroider som dyker upp. Asteroiderna har olika former, storlekar och hastigheter. Rymdfarkosten är utrustad med en laser för att hantera de irriterande rymdstenarna. Om de avfyrade projektilerna når något av målen bryts de i mindre bitar. Ett kontrollsystem måste ta hänsyn till alla olika funktioner i systemet och bestämma rymdfarkostens rörelse- och skjutbeslut” (https://xfuzzycomp.github.io/XFC/about.html).

Kan du skapa en AI (ett kontrollsystem) för Asteroids som inte bara manövrerar skickligt utan även kan förklara sina beslut om eld och rörelse på ett lättbegripligt sätt? Detta är utmaningen i *Explainable Fuzzy Challenge* (XFC), en årlig tävling för studenter och doktorander inom förklarbar AI (XAI).

I XFC 2026 vann en KEX-grupp från denna kurs första pris för ”Most Innovative Award” och ”Best Explanations Award”. Återstår nu bara första pris för ”Competition Playoff Results” …

Syfte
Projektet syftar till att bygga en AI för spelet Asteroids som kan förklarar sitt agerande. Projektgruppen väljer själv om den vill delta i XFC 2027.

Metod
Se https://xfuzzycomp.github.io/XFC/rules.html. Projektgruppen kan fritt välja vad för slags AI-metod som ska användas.

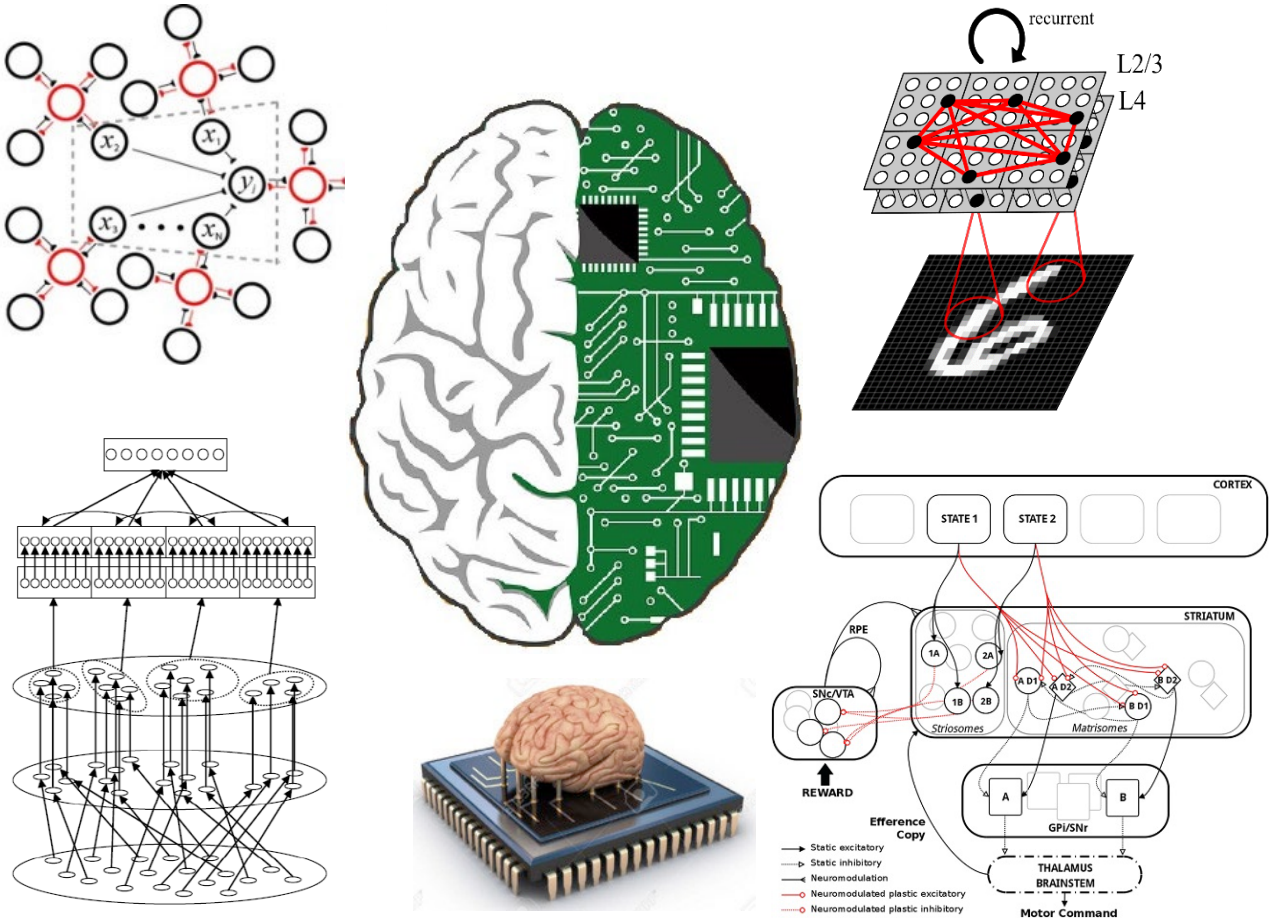
Kontext
Projektet bedrivs i anslutning till forskningsprojekt om taktisk AI på FOI.

Introduktion till XFC
Se https://xfuzzycomp.github.io/XFC/seminars.html., specifikt följande videos:

- Introduction to Asteroid Smashers, https://youtu.be/n9nf_EpCzck
- Introduction to Fuzzy Logic, https://youtu.be/I3kN0NRIyyA
- Free Optimization, https://youtu.be/6qfPmJHaBNE

Context O: Computational brain modelling and brain-like computing

Context Responsible: Pawel Herman (paherman@kth.se)
Division of Computational Science and Technology (CST)



The general focus is here on developing, studying and/or applying connectionist (neural network based) models of the brain. The proposed topics range from simulating detailed spiking neural networks to investigating and validating more abstract brain-like computing architectures. Projects can be formulated to either address theoretical questions, develop neuromorphic methods or test the networks' functionality in applications.

Please bear in mind that project details and specific research questions within the proposed themes are discussed individually with students depending on their interests. There is also a lot of flexibility in defining the scope and size of these projects. Some project ideas at the cross-sections of the following themes can be proposed/found. Students will have an opportunity to learn to use dedicated simulation software (with a possibility to rely on Python interface) or exploit their programming competence to build their own computational tools for theoretical or applied research. The focus however is on the scientific essence of the project, not on the methodology used.

The suggested projects are organized in two main themes, each of which describes a set of proposed topics. The lists of topics and some project ideas are not meant to be limiting in any sense and can therefore be easily expanded by students' own ideas.

Project O1: Simulations of attractor neural networks as models for human memory

Supervisor: Pawel Herman (paherman@kth.se)
Division of Computational Science and Technology (CST)

General theme

There have been a range of theoretical concepts of brain computations proposed in computational neuroscience. Among the connectionist (network-based) approaches to modelling brain function, an attractor theory of neural computations has recently received particular attention. The functionality of attractor networks has been found helpful in explaining various perceptual and memory phenomena. Consequently, these models can be considered as fundamental components of systems level approach to modelling brain function within the framework of network-of-networks architecture. An implementation of attractor memory models can range from a more biologically plausible networks of spiking neurons to more abstract networks of units with continuous rate-based input/output.

More biologically detailed models with spiking neurons and synapses provide an opportunity to study rich neural dynamics in close relation to biological data, and specifically, recordings from the brain tissue. This way both dynamical and functional aspects of fascinating cortical phenomena can be studied. Such spiking neural network models are usually developed using dedicated simulation software, e.g. Nest, Neuron, Genesis etc.

More abstract networks relying on rate-based units (i.e. with non-spiking real-valued input/output like in more classical artificial neural networks) on the other hand allow for constructing larger systems with the aim of exploring functional aspects of the simulated attractor memory system. In this context, both generic theoretical investigations into computational capabilities of memory (learning, memory capacity etc.) as well as specific applications in pattern recognition, whether in a biological or non-biological data mining context, can be pursued.

Within this theme other computational theories of the brain, e.g. liquid (echo) state machines or other recurrent architectures, can also be studied. In this regard, computational or dynamical aspects as well as application-oriented questions may be explored. Students can make use of existing software simulators or developed their own implementations of network models.

Project ideas

1. Studying the effect of different connectivity patterns, levels of modularity, network architectures, and their dimensionality on the dynamics and function of the attractor or other network models of selected brain function (memory, perception etc.).
2. Investigating the effect of different local learning rules on functional and operational capabilities of network models of the brain's memory function.
3. Exploring capabilities of attractor-like network memory models to operate on sequential memories (sequence learning capabilities, sequential mental planning etc.).
4. Investigating the storage capacity of brain's memory models for long-term memories acquired over the lifespan – simulation of long-term memory maintenance in our brains throughout our lives (subject to varying age-dependent processes).
5. Investigating the sensitivity of the model to the level of biological detail being accounted for (discussion on the required level of complexity and the relevance of biological constraints).

Project O2*:* Brain-like neural networks and associated neuromorphic computing al.gorithms – theoretical developments and applications

Supervisor: Pawel Herman (paherman@kth.se)
Division of Computational Science and Technology (CST)

General theme

Development of brain models to study neural phenomena, as broadly discussed in the first theme above, often leads to better understanding of the nature and purpose of neural computations. From a broader perspective, these computations can be seen as an inspiration for novel approaches to generic information processing. Good reputation of neural network architectures in this regard is largely due to the impressive capabilities of information processing in the brain, which robustly handles large volumes of noisy multi-modal data received in continuous streams. Consequently, brain-like computing has long been considered as a particularly appealing concept in a broad field of information and, nowadays, data science. With the increasing availability of powerful computing platforms and intensive development of neural network based brain models as well as a growing body of knowledge about computational mechanisms underlying brain function, there is a surge of interest in adapting these functional aspects to devise algorithms in the form of brain-like neural networks for more generic applications in the field of data mining, pattern recognition etc. These efforts are urgently needed and particularly relevant to real-world problems involving so-called big data, for example in exploratory analysis of large volumes of high-dimensional neuroimaging data for research or clinical purposes.

Brain-like neural networks differ from today's deep neural networks in the connectivity patterns, architectural design and learning methods among others. In consequence, they can help addressing different class of problems, e.g., in unsupervised or semi-supervised scenarios with lower demand for data and stronger demands for robustness. They can be deployed as spiking neural networks (event based asynchronous systems) in neuromorphic applications or as more abstract rate based neural network models (synchronous with continuous output similar to traditional neural networks) depending on the problem, specific application needs or given other constraints.

Project ideas

1. Adapting selected brain-like computing paradigms for large-scale data mining, e.g. to perform exploratory search for patterns in data.
2. Devising new brain network inspired approaches to generically process temporal or sequential data and/or comparing to the existing state-of-the art attempts.
3. General evaluation, benchmarking and validation of brain-like computing algorithms on different pattern recognition problems.
4. Testing robustness (sensitivity analysis, noise handling capabilities, computational speed) and benchmarking brain-like computing methods against more conventional machine (/statistical) learning techniques on a selected set of benchmark problems.
5. Exploring capabilities of recurrent brain-like neural networks for more challenging pattern recognition tasks corresponding to human perception, e.g. pattern completion, gestalt perception, clustering, denoising, figure-ground segmentation etc.

Context P: Sign Language Processing

Context Responsible: Jonas Beskow (beskow@kth.se)
Division of Speech, Music and Hearing





In recent years, we have witnessed a revolution in the development of machine learning models capable of processing natural language. Large language models (LLMs) have shown remarkable capabilities of understanding, producing and translating language. There is, however, one class of natural languages that have been largely left out of this rapid development, and that is the signed languages. There are about 200 signed languages, used by 70 million people worldwide.

Signed languages represent a special challenge since they, in contrast to spoken languages, have no universally adopted written form. For storage and transmission, one has to rely on video. For communication between signers and non-signers, either costly interpreter services are required, or one has to resort to limited text-based communication - which will be in a second language to a native signer. The promise and potential utility of sign language (SL) technology is thus substantial in terms of reducing communication barriers, allowing for signers to use language technology in their native language. Despite all of this, the progress in SL technology has been limited in comparison to the rapid development for spoken languages.

The projects deal with interesting and challenging problems in sign language processing, in many cases based on techniques already proven for text and speech. You will get the opportunity to work with large datasets, multimodal language models, and deep-learning for both supervised and self-supervised learning. For some of the projects, previous experience with deep learning frameworks is recommended, but there is flexibility in the projects with regards to the exact methods used.

Project P1: Sign Spotting

Supervisors: Jonas Beskow (beskow@kth.se), Yunus Can Bilge (ycbilge@kth.se)

Sign spotting[1] refers to identifying individual signs (a.k.a. glosses) that occur in continuous sign language (compare to *keyword spotting* in speech). A sign spotter can be the first step in a sign language translation pipeline (sign-to-text) or it can be used for other tasks such as search/retrieval. A fixed vocabulary sign spotter can be implemented as a classifier with one class per sign to be recognized. A more flexible, open vocabulary sign spotter uses a dictionary of signs, stored in some suitable representation, which are compared to the continuous sign language stream.

The goal of this project is to implement one or more sign spotting solutions and evaluate them on a standard dataset, either comparing a fixed- vs open vocabulary spotter, or compare different choice of representations (e.g 2D keypoints vs image features)

[1] Varol, G., Momeni, L., Albanie, S., Afouras, T., & Zisserman, A. (2022). Scaling up sign spotting through sign language dictionaries. International Journal of Computer Vision, 130(6), 1416-1439.

Project P2: Recognizing signing attributes using vision-language models

Supervisors: Jonas Beskow (beskow@kth.se), Yunus Can Bilge (ycbilge@kth.se)

In sign language processing, the ability to detect individual low-features of sign language (also known as phonetic properties - hand shape, location/place of articulation, movement, palm orientation) is important. Traditionally, such models require detailed and expensive frame-level labelling of sign language data. Recently, multimodal language models have gained strong capabilities when it comes to zero-shot analysis of images and video.

In this project, the task is to use state-of-the-art VLM models (e.g. Llava, Gemini, GPT-5) to analyze sign language frames in order to extract phonetic properties. The goal is to establish an efficient LLM-based pipeline for labelling. Since VLM:s can be slow/expensive to run, an important step is to pre-select which frames to analyse, e.g. based on motion dynamics in the video. These will then be evaluated on existing labelled datasets, for example Swedish SL dictionary [1]. The efficiency of the VLM labelling is then compared to the solutions based on linguistic annotation of phonological features, such as SignWriting, and models trained on this data, for example [2].

[1] Svenskt teckenspråkslexikon. (2025). Swedish Sign Language Dictionary online. Stockholm. https://teckensprakslexikon.su.se
[2] Zilberman, R., Moryossef, A., Langer, O. SignWriting Transcription. (2024). https://github.com/sign-language-processing/signwriting-transcription/tree/main

Project P3: Sign language detection

Supervisors: Jonas Beskow (beskow@kth.se), Yunus Can Bilge (ycbilge@kth.se)

In a video containing a person, how can we tell if this person is performing sign language, or if they are doing some other motions, e.g. vivid gesturing while speaking? Humans can often do this quite easily, but automatic sign language detection [1] is not trivial (compare to the corresponding problem of *speech activity detection*, that aims to distinguish speech from noises, music etc). Robust sign language detection has many uses, and is critical in harvesting large in-the-wild sign language datasets e.g. from YouTube and TV.

In this project, the task is to first collect a diverse dataset consisting of equal parts sign language (from a variety of SL datasets) and other human motion, e.g. co-speech gesture, (which can be gathered e.g. form ted-talks and similar). Then, you should train a model, using architecture of your choice, that can take a short video slice (a few seconds) and determine if it is sign language or not, and evaluate the results.

[1] Moryossef, A., Tsochantaridis, I., Aharoni, R., Ebling, S., & Narayanan, S. (2020, August). Real-time sign language detection using human pose estimation. In *European Conference on Computer Vision* (pp. 237-248). Cham: Springer International Publishing.

Project P4: Masked sign language prediction model

Supervisors: Jonas Beskow (beskow@kth.se), Yunus Can Bilge (ycbilge@kth.se)

One of the most powerful and successful approaches in self-supervised learning for natural language and spoken language, has been the MLM (masked language model), as pioneered by models like BERT [1] (for text) and Wav2vec 2.0 [2] (for speech). In a nutshell, these models are trained by taking a signal (e.g. a text sentence, or a part of a speech waveform) hiding/masking out part of the signal, and then letting the model predict what has been masked out, based on the context. Once the model is trained, on a large unlabelled dataset, it can be fine-tuned to solve a specific task with a much smaller dataset. In this project, the goal is to implement an MLM-like training paradigm for sign language and train it on a large SL dataset (e.g. YouTubeASL, 1000 hours), and then fine-tune the model on a small task (e.g isolated sign recognition).
This project involves advanced deep-learning techniques and requires previous experience of training large models.

[1] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers) (pp. 4171-4186).
[2] Baevski, A., Zhou, Y., Mohamed, A., & Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. Advances in neural information processing systems, 33, 12449-12460.

Context Q: Software Integrity

Context Responsible: Martin Monperrus, TCS (monperrus@kth.se)
Division of Speech, Music and Hearing

Research project: CHAINS (https://chains.proj.kth.se)

The general topic

Modern software is assembled, not written. A single application pulls in hundreds of third-party packages, is built on a machine nobody inspects, and is deployed as a binary that nobody can trace back to its sources. Every one of these steps is a place where code can be silently altered — by an attacker, by a misconfiguration, or by accident. **Software integrity** is the engineering discipline that answers one question: *is the software running here exactly the software that was intended, and can we prove it?*

This context groups three bachelor projects that each attack one link of that chain: the **package** you depend on, the **build** that produced your binary, and the **machine** that runs it. All three are hands-on: students write code, run measurements on real open-source software, and report quantitative results. No prior security course is required — solid programming skills (Python, Go, JavaScript, or Rust) and curiosity are enough. All three topics come from ongoing research in the ASSERT group and connect to the EU/SSF-funded CHAINS project on software supply chain security.

Common to the three projects: -

Project Q1: Sandboxing dependencies: a package manager with capabilities

Supervisor: Martin Monperrus, group, EECS/TCS — monperrus@kth.se

Background. In npm, PyPI, Maven and Cargo, *all* dependencies run with the full privileges of the application that uses them. A date-formatting library has exactly the same right to read your SSH keys and open network connections as your main program. This is why a single compromised package can lead to a worldwide incident. Operating systems solved an analogous problem decades ago with *capabilities*: a component only gets the specific rights it declares it needs.

The project. Design and prototype a capability-based dependency mechanism for one package ecosystem of your choice. Concretely: 1. Define a small capability vocabulary (e.g. read:fs, write:fs, net, exec, env) and a manifest format in which a package declares what it needs. 2. Implement enforcement for a real runtime — e.g. Node.js permissions, WebAssembly + WASI, Linux seccomp/Landlock, or Deno's permission model. 3. Take 20–50 popular real packages, infer the capabilities they actually use, and measure: how many packages need almost nothing? How much do declared and observed capabilities differ? What is the runtime overhead?

Research question. How much privilege do real-world dependencies actually need, and what is the cost of enforcing it?

Skills used/learned. Systems programming, sandboxing (WASM/seccomp), dynamic analysis, ecosystem measurement.

Starting points. - E. Wittern et al., *A look at the dynamics of the JavaScript package ecosystem* - WASI capability-based security model: https://wasi.dev - Linux Landlock: https://landlock.io

Project Q2: Can you trust a build attestation?

Supervisor: Martin Monperrus, group, EECS/TCS — monperrus@kth.se

Background. A *build attestation* is a signed statement saying "this binary was produced from this source, by this builder, at this time". GitHub, Sigstore and the SLSA framework now generate them automatically for millions of artifacts. They are the backbone of the industry response to supply chain attacks. But an attestation is only as strong as the assumptions behind it — and those assumptions are almost never written down. Who signs? Who can revoke? What happens if the build machine itself is compromised?

The project. Systematically dissect the trust model of real build attestation systems. 1. Set up 3–4 attestation pipelines end to end (GitHub Artifact Attestations, Sigstore/cosign, SLSA generators, Tekton Chains) on a toy project and on a real open-source project. 2. Write down, for each, an explicit **trust model**: what each party must be trusted for, what an attacker who controls each component can achieve. 3. Demonstrate concrete weaknesses in a controlled environment: e.g. an attestation that is valid but describes a build you can influence; a verification step that silently succeeds when there is nothing to verify. 4. Harvest attestations from public registries and measure how many are actually verifiable and how complete their provenance is.

Research question. Do deployed build attestation systems deliver the guarantees their users believe they deliver?

Skills used/learned. CI/CD engineering, digital signatures and PKI, threat modelling, empirical measurement of public registries.

Starting points. - SLSA framework: https://slsa.dev - Sigstore: https://sigstore.dev - Lamb & Zacchiroli, *Reproducible Builds: Increasing the Integrity of Software Supply Chains*, IEEE Software 2021

Project Q3: Integrity monitoring of servers with checksum databases

Supervisor: Martin Monperrus, group, EECS/TCS — monperrus@kth.se

Background. Once software is deployed, the question shifts from “was it built correctly?” to “is it still what we deployed?”. Servers accumulate unauthorized changes: a modified configuration file, an injected web shell, a silently patched binary, or plain data corruption. The classical answer is *file integrity monitoring*: hash everything, store a baseline, and re-check periodically. The classical answer is also widely known to be noisy and therefore widely ignored — which is exactly what makes it an interesting engineering problem.

The project. Build and evaluate a modern server integrity checker. 1. Implement a scanner that walks a filesystem, computes cryptographic checksums, and maintains a tamper-evident baseline database (e.g. an append-only Merkle-tree log, so the baseline itself cannot be quietly rewritten). 2. Cross-check observed files against *upstream provenance*: distribution package databases (dpkg/rpm hashes), language registries, and container image layer digests. What fraction of a real server’s files can be traced back to a known-good published artifact? 3. Deploy it on real machines (a KTH lab VM, a container image, a Raspberry Pi) over several weeks and quantify the actual signal-to-noise ratio: how many changes per day, how many are benign, what heuristics kill the noise? 4. Validate detection against injected modifications.

Research question. How much of a running server’s state can be cryptographically attributed to a trusted origin, and can integrity alerts be made quiet enough for humans to act on?

Skills used/learned. Systems programming, cryptographic hashing, Merkle trees and authenticated data structures, longitudinal experimental evaluation.

Starting points. - AIDE / Tripwire / osquery (existing file integrity monitoring tools) - Crosby & Wallach, *Efficient data structures for tamper-evident logging* - Go checksum database and transparency logs: https://sum.golang.org