



Projektvalskatalog

EF112X kandidatexamensarbete inom elektroteknik (15 hp) våren 2026

I denna katalog kan du hitta information om alla valbara projekt inom kursen EF112X kandidatexamensarbete inom elektroteknik (15 hp) som erbjuds våren 2026 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 finns 77 projekt att välja emellan inom ramen av 18 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: Social Networks: Analysis, Learning, and Influence (*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 H: Predicting the future sustainable power system (*Lars Nordström*)
- Kontext I: Simulation, Design, and Evaluation of BioMEMS (*Thomas Winkler*)

Kontext inom elektromagnetism, fusion och rymdteknik

- Kontext J: Design and testing of novel microwave/antenna technologies (*O.Q. Teruel*)
- Kontext K: Fusion solens energikälla på jorden (*Thomas Jonsson*)
- Kontext L: Solar wind and planetary environments (*Shane Carberry Mogan*)

Kontext inom information och nätverksteknik

- Kontext M: Machine learning over networks (*Carlo Fischione*)
- Kontext N: Cybersecurity (*Hamed Nemati*)
- Kontext O: Information engineering: Big Data & AI (*Tobias Oechtering*)

Kontext inom datavetenskap och maskininlärning

- Kontext P: Computational brain modelling & brain-like computing (*Pawel Herman*)
- Kontext Q: AI, games, and strategy (*Mika Cohen*)
- Kontext R: Sign language processing (*Jonas Beskow*)

Tillgängliga projekt 2026

Alla projekt som erbjuds vårterminen 2026 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 "skraddarsy" 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:** Fredag, 10 okt 2025, kl 15:15-18:00, sal B1
- **Projektval:** 1-15 nov 2025
- **Kursstart:** vecka 3, 2026 (uppstart-möte, i anslutning första träff med handledaren)
- **Kursslut:** vecka 21, 2026 (heldag KEX-presentation)

Kurs-PM

All information om kursens uppbyggnad finns i kurs-PM. Kurs-PM:et läggs ut senast den 1 november 2025 på projektvalsidan: www.kth.se/social/course/EF112X/

Behörighet

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.

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 är från en annan KTH skola (fysik, farkost, teknisk matematik eller energi och miljö programmet), 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 själva kandidatexjobsprojektet du vill jobba med.

När sker valet?

Valet av kandidatexjobbprojektet görs under perioden 1-15 november 2025. Resultatet påverkas ej av när du väljer under valperioden.

Projektgruppen

Kandidatexjobsprojektet 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 mest intressanta projekten ur denna projektvalskatalog. 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 bådars 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!

Anita Kullen (kullen@kth.se)

Kursledare för kandidatexjobbskursen EF112X
Stockholm, 4 oktober 2025

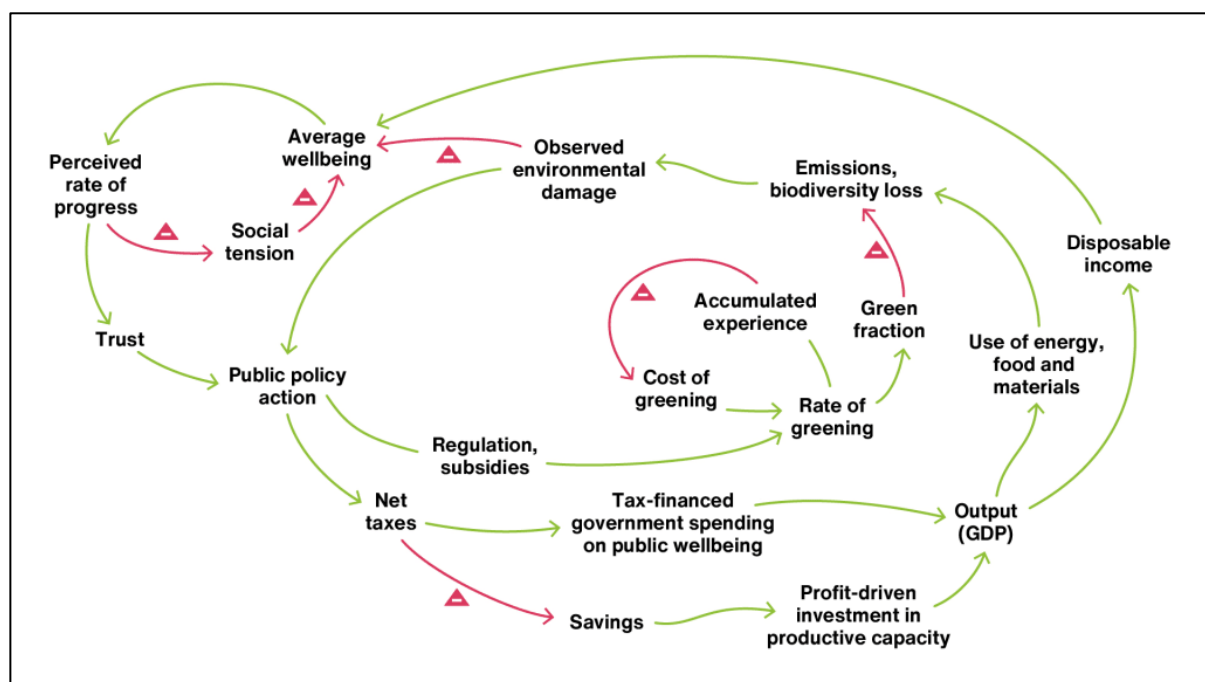
Context A: The Dynamics of a Sustainable Society

Context Responsible: Matthieu Barreau, barreau@kth.se
Divison of 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 is then taking into consideration environmental concerns along with economic development and social aspects. The 17 Sustainable Development Goals is 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, predictions 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 between 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 to 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 a 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, for example 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 is 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 have a physical sense. This model should 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.

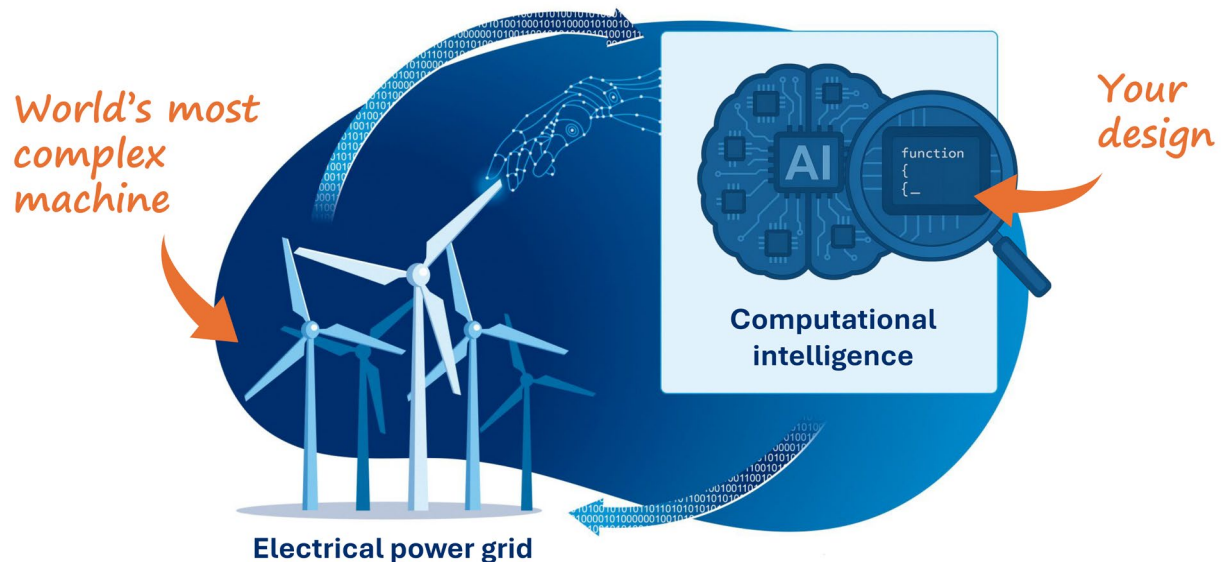
In this thesis, we propose the application of RL techniques to a subsystem of the World3 model, a well-known system dynamics model used for studying global sustainability. Using RL the students should try to derive the optimal control policy with respect to a given cost. Constraints on the measurements (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 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.

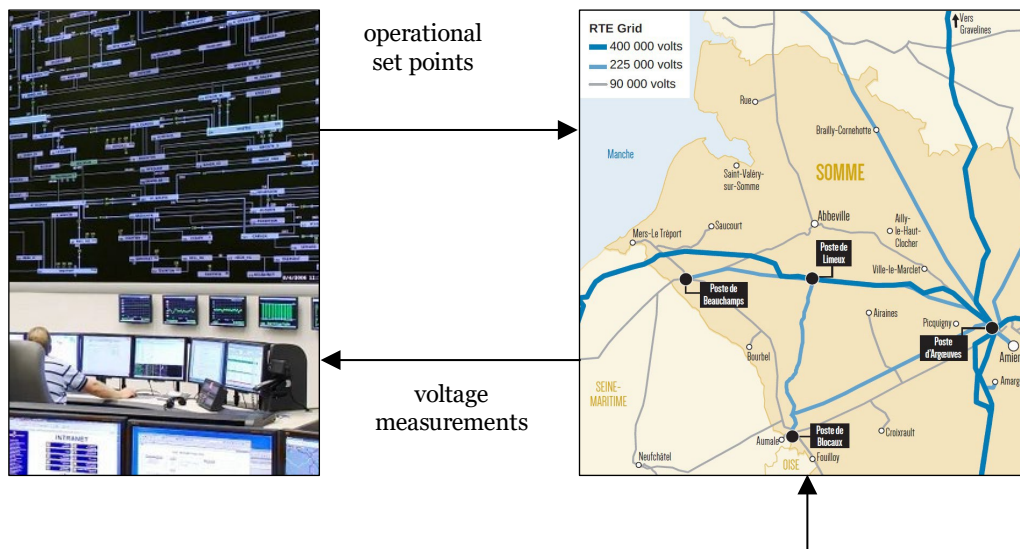


Figure 2: In future electrical power grids, generators and other grid assets will be autonomously controlled in real-time via smart algorithms (i.e., AIs) and without any human intervention.

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: Social Network Dynamics: Analysis, Learning, Influence

Context Responsible: Angela Fontan, angfon@kth.se
Division of 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, models of opinion formation have been developed to capture how the opinions 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: How can individuals maneuver each other to gain social power (= amount of influence they have on social discussions)? How can we learn the structure of interpersonal influences (often key to predicting and analyzing behaviors in a social network) if it is unknown? What risks arise when malicious agents try to destabilize communities by spreading misinformation?

The three projects in this context address these complex challenges; they make use of both synthetic networks and real-world datasets to characterize social networks in numerical simulations to illustrate the findings.


Projects C1 – C3

This context has 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. However, you will all share common datasets and the same 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.

Project C1: Influence and Social Power in Social Networks

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

This project C1 explores how individuals' opinions evolve within a social network. Using graph theory, people are modeled as nodes and their relationships as edges, which allows for the identification of key individuals in the social network. Known models of opinion evolution are applied to examine how information or influence spreads in the social network. Simulations will explore how the “social power” of each individual evolves over time.

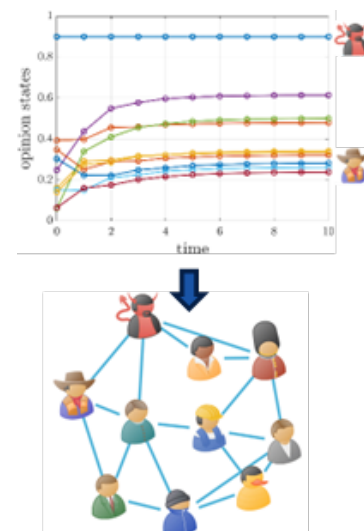


1. Study key concepts in social networks to identify the most “influential” nodes/individuals in a social network (such as centrality measures and social power);
2. Illustrate these concepts using both synthetic networks and real-world datasets, such as the Stanford Large Network Dataset Collection (SNAP, <https://snap.stanford.edu/data/>);
3. Using known models in social network dynamics, simulate and analyze how influence and social power propagate through a social network, using both synthetic and real-world networks;
4. Relate the findings from the simulations to the concepts derived at point 1.

Project C2: Learning Structures in Social Networks

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

This project C2 focuses on learning the underlying structure of a social network from observed interactions. We explore methods that infer social influences in a group of individuals whose opinions are supposed to evolve according to a prespecified (known) model.




1. Study key structural concepts in social networks (such as sparsity and density);
2. Illustrate these concepts using both synthetic networks and real-world datasets, such as the Stanford Large Network Dataset Collection (SNAP, <https://snap.stanford.edu/data/>);
3. Using known methods in social network dynamics to estimate the social interaction network and the strength of the connections, given measurements of the evolution of opinions;
4. Validate the methods through simulations on both synthetic and real-world networks.

Project C3: Malicious Agents in Social Networks

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

The project C3 investigates how the presence of malicious agents can disrupt and destabilize social communities by fostering disagreement and increasing polarization. Known models of opinion dynamics are used to study how such attacks can influence societies.



1. Study key discord concepts and measures in social networks (such as disagreement and polarization indexes);
2. Study the problem of maximizing discord from the perspective of a malicious agent;
3. Illustrate these concepts using both synthetic networks and real-world datasets, such as the Stanford Large Network Dataset Collection (SNAP, <https://snap.stanford.edu/data/>);
4. Explore how a social planner (the defender) can implement an intervention strategy to minimize the discord that the malicious agent is trying to maximize.

Context D: Electric Transportation

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

Introduktion

Dagens samhälle är starkt beroende av ett fungerande transportsystem som ska klara av såväl människor som gods. På global nivå står transporterna 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*; det behövs elektriska farkoster för såväl land, luft och vatten.

I kontext D studeras flera farkoster som har det gemensamma att de ska utföra en uppgift så effektivt som möjligt. Det är en hyperavancerad racingbil, en solcellsdriven drönare, en laddbar elbåt samt flottan av laddbara personbilar. I en del av projekten ingår att bygga och utvärdera utrustning som kan effektivisera deras prestanda.

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. A laboratory mock-up should be built to demonstrate one or more of the issues as a way to communicate the research project results

Project D2: Grid reinforcement versus flexibility from electric vehicles: a case study on a Swedish distribution system

Supervisor: Dita Anggraini (ditaa@kth.se)

Background

Electric vehicle adoption has grown significantly in recent years and is expected to keep rising. It is estimated that 3.8 million EVs would have an aggregated battery capacity of 114 GWh in Sweden, which theoretically can supply electricity for Sweden for several hours [1]. If not managed properly, the growing numbers of EVs can introduce peak loads that can cause congestion in the distribution network. Traditionally, such issues are addressed by reinforcing the grid (upgrading the distribution transformers, cables, or substations). Grid reinforcement is however capital-intensive and takes a long planning horizon.

And alternative solutions are to use the flexibility from demand-side resources, such as EVs by controlling the EV charging. These resources can help to reduce or shift peak loads, potentially postpone or even avoid the needs for grid reinforcement.

Objectives

This project aims to compare the costs and benefits of traditional grid reinforcement versus using flexibility from EVs in a selected Swedish distribution system.

The following steps are recommended:

Literature review: review challenges of electrification in distribution system and flexibility concepts from EVs. Identify methods to estimate the cost of grid reinforcement.

Cost estimation of grid reinforcement: after selecting the test system and collecting relevant data (load profiles, grid capacity limits, grid component costs, etc.), the costs for upgrading the grid can be estimated.

Modelling flexibility: define the scenarios for EV smart charging and quantify the extent to which flexibility can reduce grid congestion.

Results: compare both cases and analyze the cost and effectiveness of each approach. Summarize findings and write a bachelor thesis report.

References

[1] Power Circle, "What is V2G - Vehicle to Grid?," Power Circle, Sweden, Jan. 2024. [Online]. Available: <https://powercircle.org/wp-content/uploads/2024/03/240311-V2G-faktablad-uppdatering-jan-2024-Engelsk-version-1.pdf>

Project D3: Design and optimization of an isolated HV-LV DC/DC Converter

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

Division of Division of Electronics and Embedded Systems

Background

Formula Student is a global engineering competition where university students design, build and race small formula-style cars. The KTH Formula Student vehicle, DeV19, is an electric racing car featuring a custom-built low-voltage battery. The battery comprises of 6S2P lithium-ion cells, which together form a 24V, 13.2 Ah power source. Under peak performance, the battery is expected to draw up to 10 A continuously. However, this configuration introduces challenges related to capacity, extra weight and reliability.




Fig. 1. Picture of current low voltage battery.

The Project

This project focuses on the design and optimization of a galvanically isolated, 600 V to 24 V DC/DC converter to power the low voltage system of the KTH formula student racing car.

Research Questions

- Which topologies and design choices are most suitable for the required capabilities?
- Which methods can be implemented to ensure a stable output voltage for varying loads and fluctuations on the high voltage side, while maximizing efficiency?

Tasks

- Research and selection of topology, control strategy, and high-frequency transformer design.
- Development of an isolated, 600 V to 24 V DC/DC converter with minimized losses, EMI and weight.
- Construction of the converter and functional verification in all operating points.

Project D4: Optimisation of a UAV solar power system

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

Background

Integration of solar power in a UAV (Unmanned Aerial Vehicle) system promises significant increase of the flight duration, up to perpetual flight. A project at KTH focuses a dedicated UAV is being custom designed for forest monitoring and fire detection. The objective is being able to fly during the daytime (i.e. when illuminated by the sun).

The project

Integration of the solar cells with the electrical power system of the UAV requires addressing a number of issues. To harvest maximum power from the solar panels with varying illumination conditions, a dynamically adjusted maximum power point tracker (MPPT) is needed. It should also be seamlessly integrated with the battery management system (charging/discharging). Integrated measurement of currents/voltages allows for a more advanced and intelligent power handling system. Considering that the system is to be implemented on a flying platform, the solution should be mechanically and thermally robust, and lightweight. Several prototype systems have been developed within a BSc and MSc theses during 2024-25, which now need to be optimized for performance and integration with the UAV.

Tasks

The tasks in this project include:

- Getting familiar with the prototype systems
- Optimising the design of the system
- Implementing the design in a PCB
- Conducting the tests of the design

Context E: Power systems control

Context Responsible: Qianwen Xu, qianwenx@kth.se
Division of Electric Power and Energy Systems

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, effektlöden 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

Supervisors: Qianwen Xu, qianwenx@kth.se, Fei Liu fei7@kth.se Electric Power and Energy Systems

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.

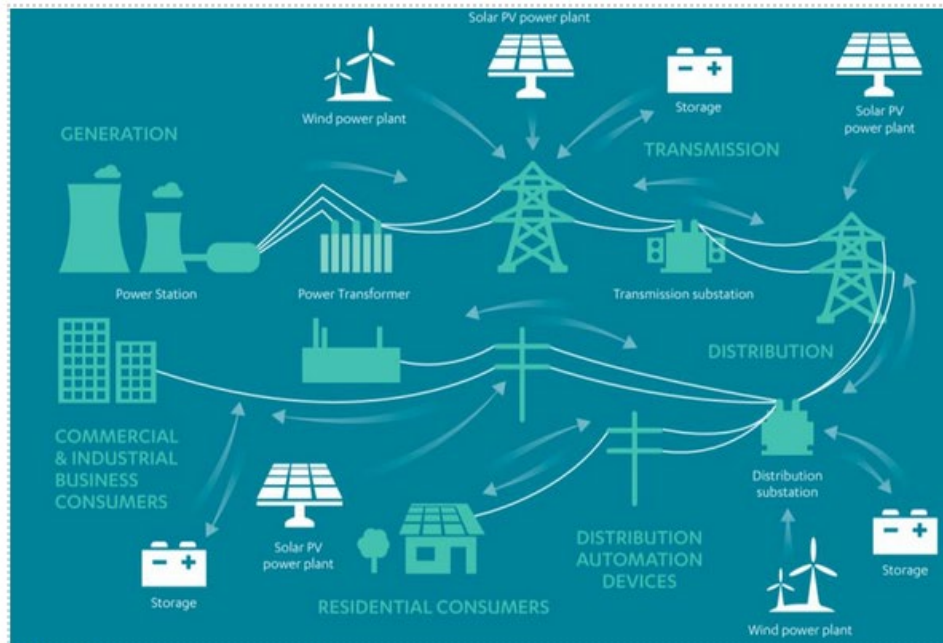


Fig: Future power system

Project E2: Frequency Control of Renewable Energy Integrated Power System

Supervisors: Qianwen Xu, qianwenx@kth.se, Xuan Jiang xuanj@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, Kamil Swiderski kamilsw@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: Real-time power system measurement

Supervisor: Valgerdur Jonsdottir, jonsdo@kth.se

A critical part of power system control is of course accurate measurements. In a large electric power system thousands of measurement points are monitored and used locally in protection and automation systems and centrally in control room applications for system-wide control. One specific kind of measurement collected at a high sampling rate are *phasors* of voltage and/or current at critical locations throughout the system. The phasor data (represented by a polar complex number [Equation]in analysis) consists of magnitude and phase-angle of the measured quantity. Measuring the magnitude of a quantity is straightforward, but measuring the phase angle requires a global reference to determine the angular separation between a measurement in two locations of the power grid. Specialised Phasor Measurement Units (PMUs) utilise time-signals from Global Navigational Satellite Systems (GNSS) - such as for example GPS to establish such a reference. By using phasor measurements from several locations in the power grid, it is possible to both measure power flows but also to analyse transient behaviour of the power system.

The project involves setting up a phasor measurement system that measures voltage locally at the lab at KTH, and one other location in Stockholm creating also local visualisations of the measurements. In addition, the PMU system shall be connected to the Nordic university network of PMUs and utilised to perform basic analysis of the Nordic power system. The project is a continuation of earlier work where local PMUs have been constructed using the OpenPMU platform - an Arduino based framework. Depending on stability of the earlier PMUs, the project can either involve further development of measurement devices, a refinement of the measurement algorithm, or use of the data for power system analysis:

- Device development: Previous work observed phase offsets between two PMUs at the KTH lab, even though they were connected to the same power outlet. The offsets were suspected to be related to the GPS signal strength. Further investigation could confirm the root cause and test improvements, with the aim of removing these offsets.
- Algorithm refinement: The current measurement algorithm is based on a least-squares curve fitting method. A possible refinement would be to implement DFT-based estimation methods and compare their accuracy against the least-squares approach. Validation could follow the compliance requirements of IEC/IEEE Std. 60255-118, focusing on steady-state frequency offsets around the nominal frequency. Performance would be reported in terms of the frequency error and the rate-of-change-of-frequency error. In addition, benchmarking the local estimates against the Nordic PMU network is of interest.
- Data analysis: Use of the data for power system analysis could, for example, include implementing an event detection module that automatically identifies disturbances and stores the corresponding raw data for later offline analysis. This would make it possible to create catalogues of events and provide a basis for further analysis. Currently, there exists only the possibility to get notifications when significant deviations occur, but no automatic raw data save function for offline analysis has been integrated.

References

Open PMU project <https://sites.google.com/site/openpmu/>
Nordic PMU data network <http://skrymer.iea.lth.se/status/>
OpenPDC project <https://github.com/GridProtectionAlliance/openPDC?tab=readme-ov-file>
M. K. Penshanwar, M. Gavande and M. F. A. R. Satarkar, "Phasor Measurement unit technology and its applications - a review," *2015 International Conference on Energy Systems and Applications*, Pune, India, 2015, pp. 318-323, doi: 10.1109/ICESA.2015.7503363.
A. Phadke and J. Thorp, *Synchronized Phasor Measurements and Their Applications* (Power Electronics and Power Systems). Boston, MA: Springer US, 2008.
"IEEE/IEC International Standard - Measuring relays and protection equipment - Part 118-1: Synchrophasor for power systems - Measurements," IEC/IEEE 60255-118-1:2018, pp. 1–78, Dec. 2018.

Context F: Hydro power modelling

Context Responsible: Mikael Amelin, amelin@kth.se
Division of Electric Power and Energy Systems



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 very 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 a large

challenge to build good hydro power models for European energy system models based on public data. Therefore, there is a need to test and verify different methods for data collection of for example inflow.

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]. Software for collecting data is under development in ongoing research projects and will eventually be integrated in the Spine environment.

Project F1: Open-source model of a large Swedish river system

Supervisor: Mikael Amelin, amelin@kth.se, Electric Power and Energy Systems

The aim of this project is to set up a model of one of the large Swedish hydro power rivers in the open-source software Spine and to run simulations to study the impact of different environmental conditions. A specific objective of the project is to set up the model to easily import data from public sources and to run the simulations for longer time periods than in earlier studies.

Project F2: Balancing reservoirs

Supervisor: Mikael Amelin, amelin@kth.se, Electric Power and Energy Systems

A hydro power plant typically uses water from a reservoir to power the turbines. The water from the turbines is released into the downstream river and will eventually reach the next hydro power plant in the river. The idea of this project is to study the possible benefits of adding a regulation reservoir directly downstream a hydro power plant, which would allow the flow in the river to be controlled more independently from the generation in the hydro power plant. The idea will be tested either on a small test system or an existing Swedish river.

Project F3: Comparison of open-access hydro power databases

Supervisor: Mikael Amelin, amelin@kth.se, Electric Power and Energy Systems

Time series of inflow and run-of-the-river generation (i.e., hydro power generation from power plants without reservoirs, which means that the hydro generation becomes weather dependent) can be obtained from various open hydro power databases. These databases are created using different methodologies and raw data sources. Therefore, it is important to understand the differences between these databases to assist energy system modellers in their selection process. This project aims to model hydropower reservoir inflow and run-of-river generation using a provided open-source tool for selected countries and to compare the results with those of existing databases, including data acquisition. This comparison should be analysed in terms of the raw meteorological data source, methodology, results and future possibilities.

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

Context G: HVDC Supergrids for Offshore Wind

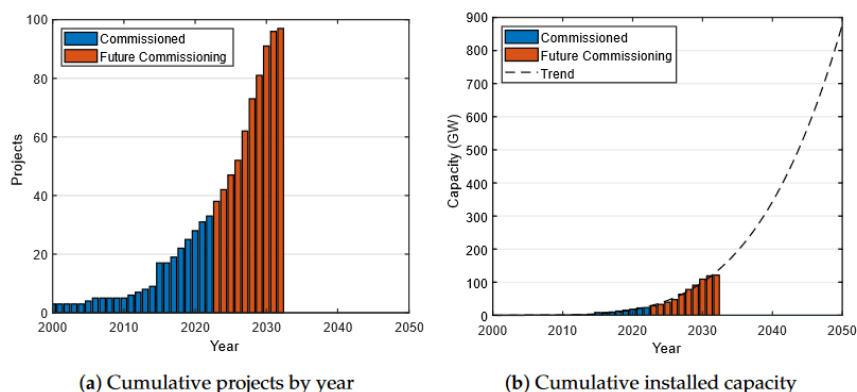
Context Responsible: Ilka Jahn, ilka@kth.se, Division of Electric Power and Energy Systems

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 and its underlying technology “HVDC”.



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

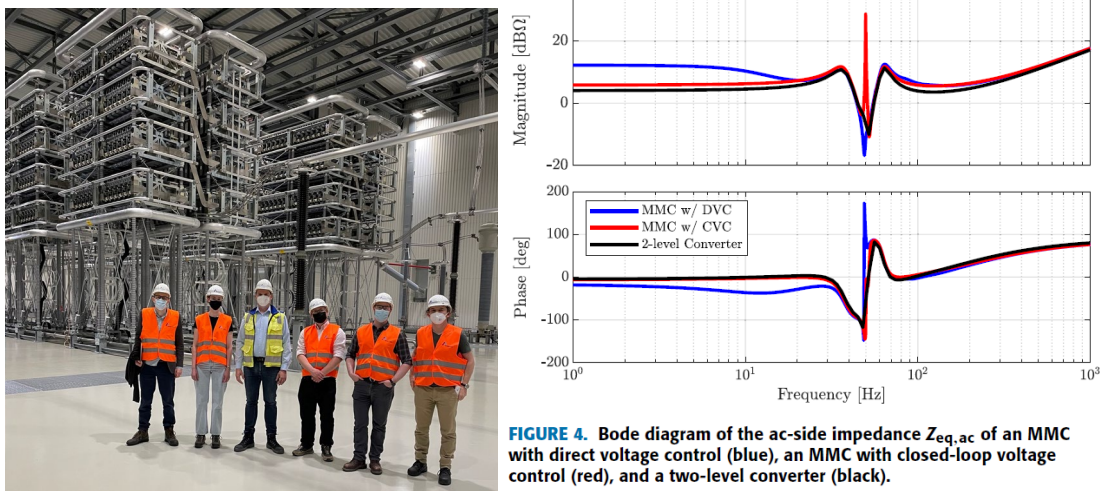


Ongoing and expected HVDC boom [2]

Project G1: Impedance-based stability of an HVDC converter

Supervisors: PhD student and Ilka Jahn, (ilka@kth.se), Electric Power and Energy Systems

Every HVDC system is built around HVDC converters that represent an equivalent impedance towards the grid. Just as in circuit theory with L and C and R elements, an HVDC system (depending on the surrounding grid) can result in a resonance point and thus unstable behavior. However, the converter impedance can be shaped with the HVDC control. In this project, we will investigate how to represent an HVDC converter as an impedance, how to scan this impedance, and how to change the impedance to mitigate unstable situations.



Picture of an HVDC converter (left) and equivalent impedance representation as Bode plot [3] (right)


In this project, the students will learn:

- About equivalent impedances of HVDC converters and how the control shapes the impedances
- About stability
- How to simulate HVDC converters in PSCAD or Matlab/Simulink (based on existing simulation) and how to scan the HVDC converter impedance in simulation

Project G2: Protection of HVDC supergrids against blackout

Supervisors: Zixuan Sui, zsui@kth.se (and Ilka Jahn), 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 anker 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 grids. This includes the detection of faults with double-ended algorithms (with information from two cable ends) requiring a communication link.



Example cable fault in an HVDC supergrid

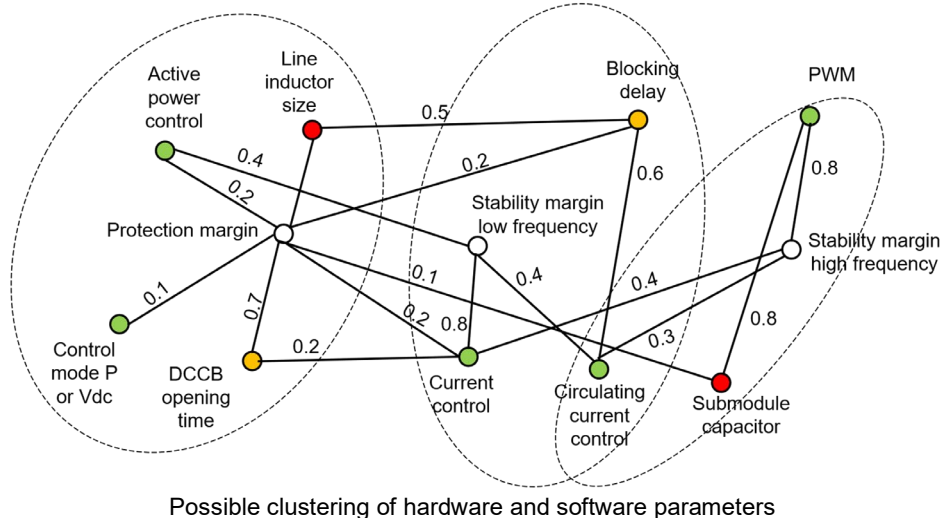
In this project, the students will learn:

- 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
- About double-ended fault detection algorithms and communication

Project G3: Control and protection interaction in HVDC

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

In traditional AC power systems, the control (in normal operation) and protection (in fault situation) have been treated separately because they act on different time-scales and the equipment is largely passive (“hardware-defined”). However, modern HVDC systems are largely “software-defined” with very fast controls and protections acting on the same time-scale, possibly disturbing the proper functionality of the other. In this project, we will investigate where and how much HVDC controls and protections interact. This includes DC-side system protection, DC-side control, and converter internal protection.



In this project, the students will learn:

- About control and protection in HVDC systems
- How to simulate HVDC technology in PSCAD (based on existing simulation)
- How to simulate control and protection in computer simulation
- How to investigate and visualize control and protection interaction in HVDC

- [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>
- [3] M. Nahalparvari, M. Asoodar, S. Norrga and H. -P. Nee, "AC-Side Impedance-Based Stability Assessment in Grid-Forming Modular Multilevel Converters," in *IEEE Access*, vol. 12, pp. 23514-23528, 2024, doi: 10.1109/ACCESS.2024.3365053.

Context H: Predicting the Future Sustainable Power System

Context Responsible: Lars Nordström, larsno@kth.se
Division of Electric Power and Energy Systems



To help mitigate climate change, it will be necessary to significantly reduce CO₂ emissions. These efforts will in turn have a major influence on power systems and electricity markets, both because power generation is in itself a large source of CO₂ emissions, but also because electricity is necessary to facilitate eliminations of CO₂ emissions in other parts of society. In short, there will be an increased demand for CO₂-free electricity generation in the future. Two main sources of such electricity is of course wind and solar power, which is being introduced in a wide scale across most electric power systems on the planet.

One important characteristic of wind and solar power is that it is non-dispatchable, i.e. the output cannot be controlled but depends on weather. As can be observed on the electricity market recently, the volatility of electricity prices has increased due to the varying inflow of power from these sources[1]. Similarly, the non-dispatchable nature of the renewable sources are creating further challenges for stable operation of the power system in real-time. To improve functioning of the electricity markets and facilitate cost-efficient and reliable planning and operation of electric power systems, there is a need for better prediction of the impact of renewable sources on the power grid. This includes several aspects of this problem including forecasting electricity prices as well as renewable generation but also load, which is also trending towards increased volatility due to new types of consumers e.g. electric vehicles, electrolyzers and battery storage systems.

The aim of this context is to explore data science based methods for improved forecasting of electricity price (F1), Renewable generation (F2) and Load (F3)

[1] <https://www.di.se/nyheter/svenska-kraftnat-varnar-for-mer-volatila-elpriser-och-okad-risk-for-akut-effektbrist/>

Project H1: 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-forecasting 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).

Project H2: Estimation and forecasting of the renewable generation

Supervisors: Lars Nordström, larsno@kth.se, and Valgerdur Jonsdottir, jonsdo@kth.se
Electric Power and Energy Systems

Wind generation in Sweden has gone through an enormous expansion the last 5-10 years and is expected to see even larger expansion with the growth of off-shore wind power [1]. Presently (2022) the total wind energy production amounted to 33TWh, approximately 20% of the total Swedish electric energy generation. Similarly, Solar power is seeing a similar expansion, albeit from lower numbers, both in terms of PV-farms in the MW scale to household level PV on roof-tops. With the size and scale of PV being smaller, and more distributed in the grid, the observability of PV is lower than that of wind generation but can during situations of low load, e.g. a warm day in July, still amount to a large proportion of the total generation.

Given the variability in in-feed from renewable sources both in time and space, there is a need to forecast the production so that planning of the operation of the power grid can be facilitated. As an example, upcoming changes in production from wind may necessitate activation of reserves and systems services. Similarly, the location of generation in the grid may cause congestion if the production is concentrated to specific areas far from load-centers. Overall, the need to predict the production from renewable generation both in space (price area) and time (hours and days) is growing.

At this background, this project has the following aims:

- (1) A review of different recent techniques developed and suggested in the literature for renewable generation (wind, solar or both) estimation and forecasting; In this review, the strengths and weaknesses of these reviewed techniques should be clarified. <https://transparency.entsoe.eu> <https://transparency.entsoe.eu>
- (2) Develop one forecasting application for renewable generation based on suitable approach identified above, using data from the ENTSO-E Transparency portal[1] applied to one or several price areas in the Nordic power system. For your application, you are encouraged to use: Python or Matlab.
- (3) To apply the developed forecasting technique to estimate-and-forecast the renewable generation. The forecasts shall be benchmarked with the forecasts available on the ENTSO-E transparency platform.
- (4) To interpret and explain the estimation and forecasting results that you have obtained

References:

- [1] <https://www.energimyndigheten.se/statistik/den-officiella-statistiken/statistikprodukter/vindkraftsstatistik/>
- [2] <https://transparency.entsoe.eu>

Project H3: Estimation and forecasting of load

Supervisors: Lars Nordström, larsno@kth.se, and Arvid Rolander arvidro@kth.se
Electric Power and Energy Systems

The electric load has in Sweden for a long period remained relatively stable. The share of residential load remaining constant with some changes in industrial and commercial load as society has moved from heavy industry in 1980s to a service based economy centered on larger cities[1]. Present forecasts [2] indicate a doubling of the electricity load, mainly due to electrification of heavy industries in mining and steel manufacturing, but other sectors such as transportation are also contributing to this growth.

Similar to the development within renewables (see project 2) the changes in load happen both in time and space. E.g. data centers are built outside municipalities or steel mills are redesigned to use electricity and Hydrogen instead of fossil fuels. Given this, there is a need to forecast the load so that planning of the operation of the power grid can be facilitated. As an example, the location of loads in the grid may cause congestion if the production is concentrated to specific areas far from these load-centers. Overall, the need to predict the loads both in space (price area) and time (hours and days) is growing.

At this background, this project has the following aims:

- (1) A review of different recent techniques developed and suggested in the literature for electricity load both residential, commercial and industrial; In this review, the strengths and weaknesses of these reviewed techniques should be clarified.
- (2) Develop one forecasting application based on suitable approach identified above, using data from the ENTSO-E Transparency portal[1] applied to one or several price areas in the Nordic power system. For your application, you are encouraged to use: Python or Matlab.
- (3) To apply the developed forecasting technique to estimate-and-forecast the total electricity load. The forecasts shall be benchmarked with the forecasts available on the ENTSO-E transparency platform.
- (4) To interpret and explain the estimation and forecasting results that you have obtained

References:

- [1] <https://www.energimyndigheten.se/statistik/den-officiella-statistiken/statistikprodukter/manatlig-elstatistik-och-byten-av-elleverantor/>
- [2] <https://www.svk.se/siteassets/om-oss/rapporter/2021/langsiktig-marknadsanalys-2021.pdf>
- [3] <https://transparency.entsoe.eu>

Context I: Simulation, Design, and Evaluation of BioMEMS

Context Responsible: Thomas E Winkler, winklert@kth.se
Division: Micro- and Nanosystems

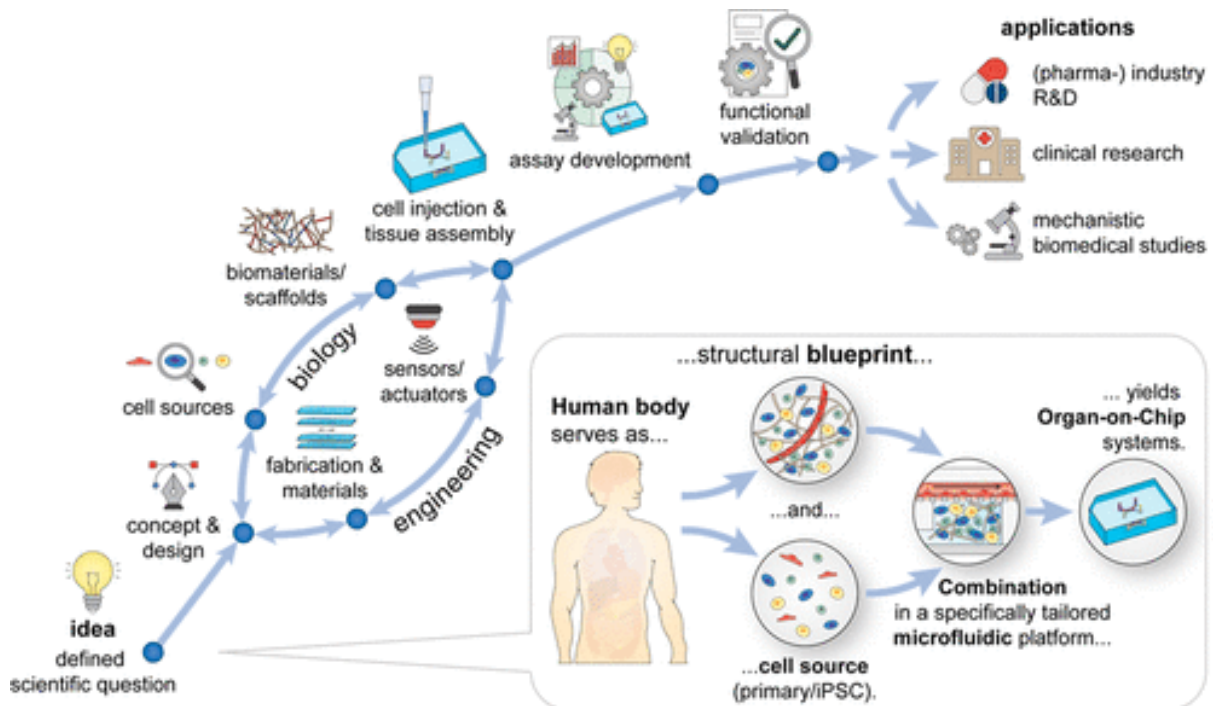


Figure 1: Organ-on-Chip Idea-to-Impact, with the focus of the available Projects highlighted in green. Credit: J. Rogal et al., ACS Biomater. Sci. Eng. (2022)

BioMEMS (bio-micro-electro-mechanical systems) are microfabricated systems that interface with living matter – where “hard” sciences and engineering meet biology. A flagship example is organs-on-chips: microfluidic platforms designed and built around the (micro)structural blueprint of the human body to create human-relevant models of health and disease (see Figure above). They are a critical puzzle piece towards Good Health and Well-being, as current models used in research and drug development are actually quite poor – animals differ fundamentally from humans, and cells in simple petri dishes drastically over-simplify the human situation. Organs-on-chips, with their engineered human-like complexity, can reduce and replace animal testing, lower rates of pharmaceutical development failures and thus costs, and shed more light on our fundamental understanding of our own cells and bodies.

This context invites you to take a systems science/engineering approach to make multidisciplinary complexity tractable, and break it down into classic aspects of the R&D loop: model and simulate to expose limits, design to balance physical and practical constraints, fabricate and implement to test hypotheses, evaluate with quantitative benchmarks – iteratively. The three projects specifically apply this to three current and quite distinct challenges we are facing in our lab regarding our organ-on-chip engineering, concerning microfabrication of membranes, integration of on-line sensors, and stabilization of liquid flows. Yet the underlying habits of thought you will practice generalize well beyond BioMEMS, powering progress from aerospace, to medical robotics, to climate tech.

No prior biology required; bring curiosity, rigor, and a scientist’s/engineer’s mindset!

Project I1: Design Challenge – How to create an invisible porous membrane?

Supervisor: Thomas E Winkler, winklert@kth.se

Thin plastic membranes with micro-pores are critical components of organs-on-chips, because they provide support for the formation of cohesive cell layers that mimic, e.g., blood vessel walls or intestinal lining while at the same time allowing these cells to interact with the surrounding environment both above and underneath (nutrients, other cells, etc.). Although cheap to produce and easy to use, standard membranes from thermoplastics like PET or PC suffer from high optical scattering as porosity increases, thus degrading the quality of microscopic imaging, typically a primary readout of organs-on-chips (see Figure below). We have identified a material/process that could theoretically overcome this limitation – and your task is to research how it could practically be implemented:

1. Learn about a range of microfabrication methods and materials, with a specific focus on membranes and their associated practical and physical limitations.
2. Starting from an existing theoretical idea, research and develop potential practical pathways for material sourcing and fabrication that would yield invisible porous membranes.
3. Critically evaluate and compare your potential strategies with regard to available resources, expected outcomes, and overall chance of success.
4. Assuming design targets can be met, implement the process/material in practice.




Figure: Similar cells imaged on an optically-clear surface (left) and on a standard porous membrane (right). Credit: M.Y. Kim et al., J. Membr. Sci. (2014)

Project I2: Integration Challenge – Electrical & fluidic interconnects for on-line sensing

Supervisor: Thomas E Winkler, winklert@kth.se

Although microscopy can reveal much about what is going on inside organs-on-chips, it is not usually feasible to keep these systems continuously under a microscope, and there are limits to what can be optically observed in the first place. For “seeing” small molecules, one very attractive approach is electrochemistry, not least because both sensors and instrumentation can be miniaturized more easily than optics. We previously used a commercial electrochemical sensor to read out metabolic information (i.e., glucose and lactate) from biological samples, but only “off-chip” (see Figure below). However, with the sensor package being in fact microfluidic, it is promising for “on-line” integration with organs-on-chips – and your task is to practically implement this:

1. Learn about electrochemical sensors and microfluidic interface rules.
2. Devise electrical (custom PCB) and fluidic components (from off-the-shelf plastics) to (a) hold two sensors in place, (b) connect to them electrically, and (c) facilitate a plug-and-play connection to the organ-on-chip.
3. Critically and iteratively evaluate possible designs in theory, as well as in practice.
4. Optionally, program an API for multi-sensor readout in Python or MATLAB.




Figure: Electrochemical monitoring of biological samples using a commercial sensor that would lend itself to integration with an organ-on-chip. Credit: I. Matthiesen et al., Biosens. (2022).

Project I3: Flow Challenge – How to perfuse a microfluidic “river”?

Supervisors: Thomas E Winkler, winklert@kth.se

Microfluidics mostly deals with completely liquid-filled pipes or ducts. However, liquid-air interfaces offer interesting physical and chemical opportunities, and it can thus be valuable to create something that is more like a river or canal (i.e., open to air) instead. For instance, we could exploit such an interface for self-assembly of biomolecules as an alternative method creating membranes inside of organs-on-chips (instead of using the plastic ones mentioned earlier). We can define the position of an air-liquid interface using an effect called capillary pinning, and have previously created a chip that exploits this over a relatively large area (see Figure below). The challenge lies in creating liquid flow using a pump system without the liquid breaking confinement and “spilling over” – and your task is to solve this:

1. Learn about microfluidic flow and microfluidic perfusion.
2. Analyze the problem using physical fundamentals/calculations and/or numerical (e.g., COMSOL, ANSYS) simulations.
3. Validate solutions using existing chips in laboratory experiments.
4. Critically evaluate theoretical/experimental results and iterate as needed.




Figure: A pinned microfluidic air-liquid interface, but without active flow. Credit: T.E.W. (unpublished)

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

Context Responsible: Oscar Quevedo-Teruel, oscarqt@kth.se
Division of Electromagnetics and Plasma Physics

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:

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 J1: Design of a dielectric lens for communications applications in W-band

Supervisors: Oscar Quevedo-Teruel, oscarqt@kth.se, Núria Flores Espinosa nuriafe@kth.se, *Division of Electromagnetics and Plasma Physics*


Due to the increase in the frequency bands of new communications systems, lens antennas are becoming very popular. Dielectric lens antennas can be used to increase the directivity of conventional antennas using low-cost materials and widely available manufacturing techniques such as 3D printing. Therefore, lenses are being considered as a suitable solution for terminals to communicate with the new generation of Low-Earth-Orbit (LEO) satellite constellations or for use in 5G and 6G communications and sensing applications.

Integrated Lens Antennas (ILAs) are a very common type of dielectric lens antennas used for beam shaping. The elliptical lens and the extended hemispherical lens are two examples of ILAs found in literature for improving the directivity of a simpler feed. The basic principle behind these lenses is their capability to collimate the radiation coming from a point source located in their focal point. However, their design entails some challenges such as reducing the reflections in the lens surfaces and adapting them to be manufactured using 3D printing.

This project will be aimed to design, manufacture and measure the S-parameters for an extended hemispherical lens fed with a WR10 waveguide working at W-band (75 – 110 GHz).

As a student, you will learn:

- The fundamentals of dielectric lens antennas
- How to simulate dielectric lens antennas in CST Microwave Studio
- How to reduce reflections at the lens surfaces
- How to match the lens model for 3D printing
- How to measure of S-parameters at the W-band



Extended hemispherical lens schematic



Photo of a 3D printed dielectric lens

Project J2: Design and Characterization of Fully Metallic Metalenses for 6G Networks

Supervisor: Oscar Quevedo-Teruel, oscarqt@kth.se, Dayan Pérez-Quintana dayanpq@kth.se, *Division of Electromagnetics and Plasma Physics*

The upcoming deployment of sixth generation (6G) of wireless networks is driving the development of devices that can address the rigorous requirements of this technology. Key features include extremely high data rates, ultra-low latency, broad bandwidth, and multi-beam operation.

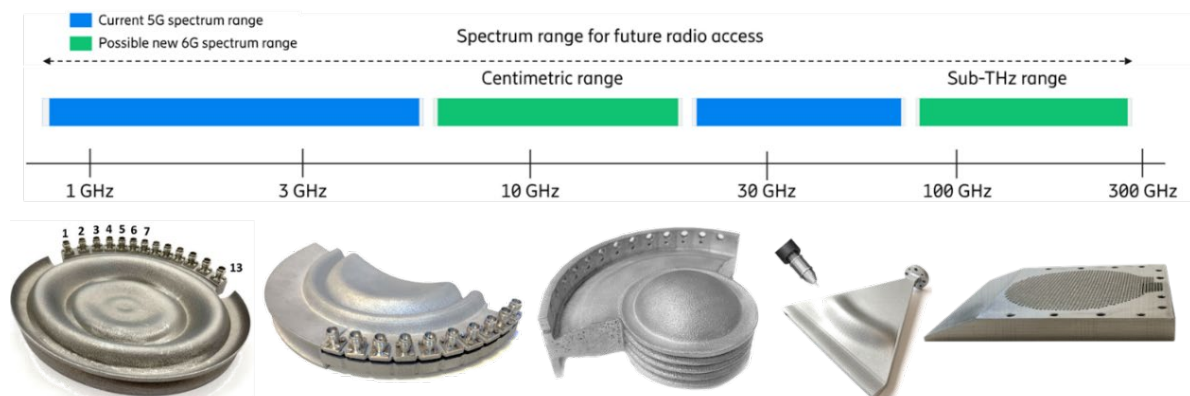
Leading industries in telecommunications, including Ericsson, have identified sub-terahertz (sub-THz) frequencies (92–300 GHz) as fundamental for 6G, while also highlighting a critical technological gap: existing antennas suffer severe efficiency losses. Higher path loss, inefficient active components, and complex feeding networks restrict current solutions. Moreover, dielectric-based technologies exhibit prohibitive losses at these frequencies, limiting the viability of conventional approaches.

As an alternative, fully metallic (FulMet) metasurfaces and flat metalenses represent a promising path. These structures can precisely control electromagnetic waves while minimizing losses. Their compact geometry, high directivity, and efficiency at millimeter waves, already exploited in 5G, make them highly attractive for extending capabilities into the sub-THz range.

This project aims to investigate, design, and optimize FulMet metalenses in the sub-THz range for 6G applications. The work will involve simulation, design, fabrication, and experimental characterization of a FulMet metalens.

As a student, you will learn:

- Fundamental electromagnetic principles of metalenses, with emphasis on Luneburg-type designs.
- How to simulate periodic structures using commercial electromagnetic software.
- How to design and implement Luneburg lenses in commercial electromagnetic software.
- How can experimental techniques be applied to measure FulMet metalenses, and what are the main fabrication technologies available for sub-THz devices.



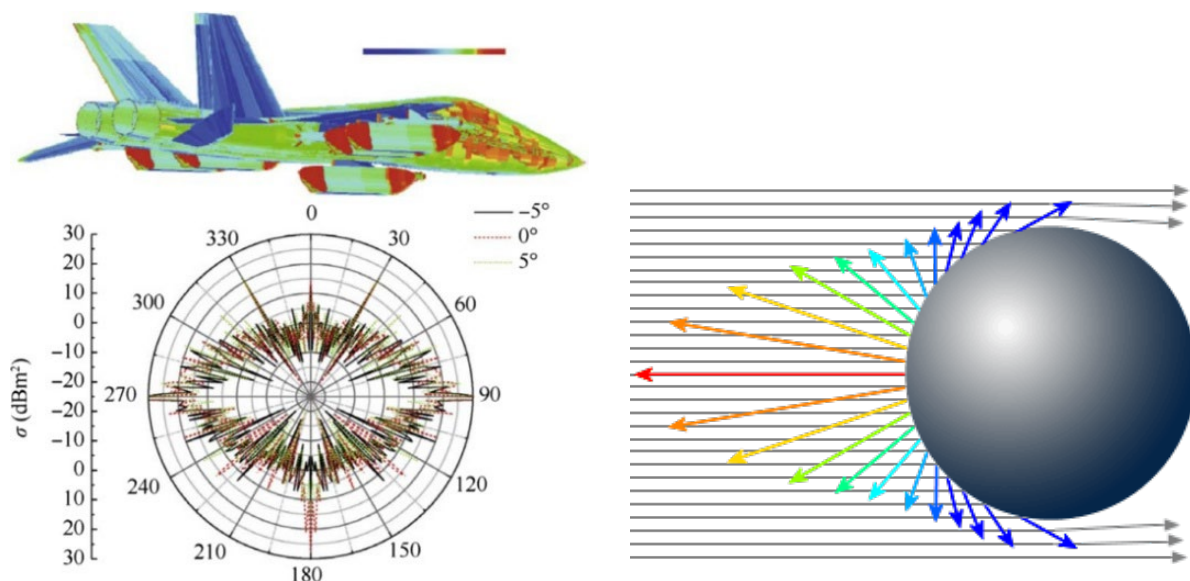
Examples of FulMet metalenses produced by CNC milling and additive manufacturing techniques.

Projekt J3: Configuration of large-scale simulations for the calculation of the Radar Cross Section (RCS) of a stealth aircraft

Supervisor: Oscar Quevedo-Teruel, oscargt@kth.se, Riccardo Musso, rmusso@kth.se
Division of Electromagnetics and Plasma Physics

The Radar Cross Section (RCS) of a structure, or radar signature, is a measure of the amount of power that is reflected and detectable by a radar. The calculation of the RCS of various objects is important for many applications, from defense to automotive radars.

In the design of a stealth aircraft, it is necessary to minimize the RCS with clever geometry and absorbing materials. However, performing a simulation of very large structures (compared to λ) is a real challenge that cannot be completed with an ordinary computer. In this project, you will learn how electromagnetic simulations work under the hood and how to set up a High Performance Computing (HPC) environment that can parallelize the calculations and drastically reduce the duration of difficult simulations. With this project, you might have the opportunity to use KTH's powerful supercomputer as development platform. <https://www.kth.se/en/eecs/nyheter/kth-s-nya-superdator-snabbast-i-sverige-1.1231754>



Left figure: radar cross section of an aircraft, shown for three cuts.

Right figure: scattering of a plane wave on a dielectric sphere.

As a student, you will learn:

- The importance of RCS for defense and communications applications and its evaluation for complex structures.
- How to go beyond the limitations of commercial software to perform massive scale simulations.
- How to configure a HPC infrastructure to parallelize and optimize scientific computing, on state-of-the-art hardware platforms.

Projekt J4: Design of a waveguide feeding network for satellite communications using electromagnetic bandgap (EBG) periodic structures

Supervisors: Oscar Quevedo-Teruel, oscarqt@kth.se, Mingzheng Chen mzchen@kth.se
Division of Electromagnetics and Plasma Physics


With the development of modern society, there is a high demand for the improvement of high-performance satellite communications. Thus, robust and efficient antenna systems in the satellite payloads are necessary, including the antenna and the feeding network. Feeding networks are guiding structures intended for power distribution in a system. They are usually manufactured in two separate metallic plates that are joint together using screws. This process is very sensitive at high frequencies resulting in misalignments and a small air gap between the two plates, producing wave leakage, decreasing the transmission and the system performance.

In recent years, periodic structures have been commonly used to design a wide range of microwave devices, such as lenses, leaky-wave antennas and flanges with low leakage. Different periodic structures possess different properties and a proper selection of them is critical for the specific application.


This project aims to study different holey EBG periodic structures for preventing the leakage in waveguide feeding networks. In particular, the project focuses on the glide-symmetric holey EBG and a novel three-hole EBG, aiming at comparing their performance in terms of leakage prevention and compactness. The work includes simulating, designing, and potentially measuring a waveguide feeding network.

As a student, you will learn:

- The electromagnetic operation of periodic structures.
- How to simulate periodic structures using commercial software.
- How to design and simulate waveguide feeding networks in commercial software.
- How to perform measurements of a waveguide feeding network.



Unit cells of Holey EBG periodic structures



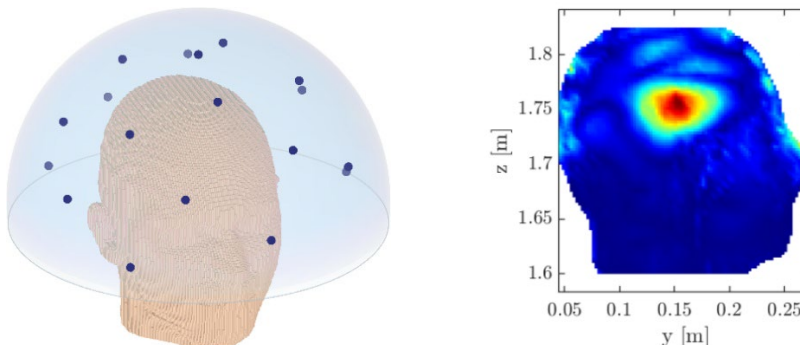
A waveguide feeding network

Projekt J5: Microwave focusing in a human brain model for deep brain stimulation

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

Stimulation of biological tissues using electromagnetic (EM) fields is of high interest within the biomedical field. One important application is neuro-modulation, used for deep brain stimulation (DBS) to treat neurological disorders like Parkinson's disease. The currently recognized treatments use intracranial electrodes with an extension wire and pulse generator to directly create intervening electric fields near neural tissues. However, this invasive approach carries risks, including surgical complications and infections. As a result, non-invasive treatments are being proposed with the prospect of achieving similar therapeutic benefits without the associated risks.

The most studied alternative methods use magnetic coils that induce low-frequency fields. Low-frequency fields can pass through the skull but cannot achieve the required focus to target a specific area of the brain. Therefore, high-frequency fields in the microwave (MW) region have been suggested to be used because of the high degree of spatial focusing. Achieving highly focused electromagnetic fields inside the human body is a complex task. Recent studies utilizing an array antenna system have shown the potential to generate such fields for DBS.



Figures: A voxel head model with antennas (the dots), and the generated EM field inside the head, obtained from work within the research group at KTH.


The objective of this student project is to obtain an understanding of how EM fields can be focused in biological tissues in the context of DBS, as well as to develop and perform numerical simulations of array antennas in voxel based human head models. The findings and insights gained from this project will be of significance for further research within the emerging field of DBS using MW fields.

In this project, the students will learn:

- About DBS and the mechanisms behind neuro-modulation.
- How to use commercial simulation software to analyze head and antenna models.
- How to evaluate EM fields created by antenna elements in an array to achieve some desired radiation pattern.
- Optimize antenna element variables for optimal EM focusing and adherence to patient safety guidelines.

Context K: Fusion, solens energikälla på jorden

Context Responsible: Tomas Jonsson (tomjons@KTH.se)
Division of Electromagnetics and Plasma Physics (EMP)



ITER - en 73m höga fusions-reaktorn som just nu byggs i södra Frankrike och som skall stå klar 2020. ITER väntas kunna producera 10 gånger mer energi än den förbrukar.

[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.

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 ($\sim 10^{20} \text{ m}^{-3}$), 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

Project K1: Estimation of relativistic electron energies in reactor-scale fusion devices

Supervisor: Mathias Hoppe, mhop@kth.se, EMF, Lorenzo Votta, votta@kth.se, EMP

Future fusion reactors of the tokamak type – the currently most promising type of fusion device – faces a major challenge. If the difficult-to-control plasma which makes up the fuel of the fusion reactor suddenly experiences an instability and cools in what is known as a *disruption*, the electrons of the plasma can be accelerated to relativistic energies and become so-called *runaway electrons* (REs), which risk causing severe damage to the reactor wall and other plasma-facing components. Since runaway electrons could potentially damage a tokamak reactor beyond repair, techniques for mitigating the damage done is one of the most highly prioritized areas of research for next-generation fusion devices such as ITER and SPARC.




Figure 1: Photo of a plasma in the TCV tokamak, where the runaway electrons can be seen in orange.

The severity of damage caused by REs is determined by their energy. Since the electrons normally lose some of their energy when colliding with other particles, one must generally formulate and solve a partial differential equation called the Fokker-Planck equation in order to calculate the energy of REs. However, by employing appropriate assumptions, one can derive a simpler equation for the average RE energy from the Fokker-Planck equation, in which the REs are treated as a fluid.

The goal of this project will be to derive an equation for the average energy of an RE, and to test numerical solutions of this model against numerical solutions of the full Fokker-Planck equation. The project will involve analytical calculations, programming, and numerical calculations with the state-of-the-art plasma physics code DREAM.

The project will be executed in the following steps:

1. Survey the relevant literature about runaway electrons and the Fokker-Planck equation.
2. Derive an equation for the average energy of electrons in a plasma.
3. Implement a solver for the equation in 2, and solve the Fokker-Planck equation using the DREAM code. Determine when the assumptions made are valid.
4. Write a report and present the method and results of the project.

If time permits, the project can be extended by also estimating the rate of production of hard x-rays by the REs. Such hard x-rays are routinely measured at tokamak experiments, and the model predictions may thus be compared to experimental data.

Project K2: Study of material segregation and enrichment

Supervisor: Per Petersson, ppeter@kth.se, EMP

Despite the magnetic confinement used in tokamaks the walls facing the plasma get 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.

An important tool that we frequently use for looking at tokamakium is Time of Flight Elastic Recoil Detection Analysis (ToF-ERDA) that gives information about all elements that are present in the surface of a material. ToF-ERDA is a type ion-beam analysis that uses MeV ions from an accelerator that interacts with the sample and knocks out some atoms from the target.

The main goal in this project will be to improve of the calibration method of the spectrums that are recorded. With the improved calibration we hope to increase the reliability by with we can determine the identity of unexpected elements and use clustering methods for determining depth profiles and separate similar elements.

For testing the results both existing data as well as data produced in this project can be used.

Introductory Part

Visit to the Tandem Accelerator Laboratory of Uppsala University and introduction to accelerator-based material analysis techniques.

Main Tasks


- Importing data from existing tools.
- Setting up several models for calibration of the data, both existing and new.
- Analysis of experimental data using the developed calibrations.
- Comparison of different methods for advances and weaknesses.
- Writing report

The main part of the work will consist of data handling using e.g. python and analysis of the results but there will also some laboratory visits and work. The project is part of larger program of experimental investigation and production of material and can be adjusted in cooperation with the student.

Project K3: Modellering av radiovågsuppvärmning i ARC

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

Fusionsreaktioner kräver mycket höga temperaturer. För att producera stora mängder energi i ett fusionskraftverk krävs cirka 200 miljoner grader. Dessa förutsättningar kan bara skapas i mycket stora och därmed dyra anläggningar. Men, med hjälp av ny supraleddande teknologi hoppas man nu kunna bygga mindre och mer ekonomiska reaktorer. I USA planerar man därför att bygga fusionsexperimentet ARC som förväntas visa på denna potential. För att nå de relevanta temperaturerna kommer ARC använda uppvärmning med hjälp av radiovågor. I denna studie kommer vi att undersöka dessa vågor, samt hur de kan accelerera snabba joner, som i sin tur överför sin energi till bränslet i form av värme.



Elektrisk fältstyrka från en numerisksimulering av radiovågsuppvärmning i ett fusionsplasma.

ARC skiljer sig mot tidigare experiment genom att man använder starkare magnetfält, vilket påverkar flera egenskaper hos plasmat (den varma joniserade gas i vilken reaktionerna ska ske). Framför allt kan man skapa plasma med högre täthet än vad som använts i tidigare fusionsexperiment. Detta påverkar i sin tur de radiovågor som används för att värma plasmat, samt hur dessa vågor accelererar snabba joner. Huvudmålet med det här projektet är att undersöka hur radiovågor kan användas för att värma ARC, hur vågen absorberas och hur snabba joner accelereras. Det är t.ex. viktigt att uppvärmningen sker i centrum av plasma, samt att vågen värmer bränslejonerna och inte elektronerna. Resultaten ska jämföras med tokamaken ITER som byggs i södra Frankrike och som har studerats i detalj i tidigare arbeten.

För att studera uppvärmningen från radiovågor måste man först lösa en vågekvation och därefter lösa så kallade *kinetiska* ekvationer som beskriver accelerationen av jonerna och hur de kolliderar med elektroner och andra joner. På KTH har vi utvecklat både en vågkod, FEMIC, och en kinetisk kod, Foppler, som ska användas i detta arbete.

I detta projekt kommer man lära sig mycket fysik. Man kommer få prova på att göra ett forskningsprojekt och dessutom göra ett bidrag till forskningen om radiovågor i ARC! I projektet kommer vi att arbeta med både COMSOL Multiphysics och MATLAB.

Målen med detta projekt är:


1. Läsa relevant litteratur om fusionsplasmafysik och radiovågsuppvärmning.
2. Lära sig den grundläggande fysiken som finns beskriven i FEMIC och Foppler, samt att lära sig köra koderna och analysera resultaten.
3. Identifiera mått som beskriver kvalitén på uppvärmningen. Dessa mått ska användas för att jämföra resultaten från ARC och ITER.
4. Ta fram parametrar som är representativa för ARC och ITER. Skapa inputfiler för simuleringar med FEMIC och Foppler.
5. Göra en kvalitativ och kvantitativ analys av simuleringsresultaten som besvarar frågorna; a) Kommer radiovågsuppvärmningen i ARC fungera som planerat? B) Hur skiljer sig uppvärmningen i ARC och ITER, och varför?
6. Skriva en rapport och hålla en presentation.

Projekt K4: Modeling EXTRAP T2R Reversed Field Pinch plasma discharge

Supervisors: Per Brunsell, brunsell@kth.se, EMF

Background

- Plasmas for research on nuclear fusion energy are typically torus-shaped plasma discharges placed inside a magnetic field coil. The type of magnetic field configuration used in the EXTRAP T2R device is called a Reversed-Field Pinch (RFP).
- Some systems in nature naturally evolve toward preferred states that are independent of the initial state of the system; the evolution toward the final state is a self-organisation process. The magnetic field configuration in the RFP is one example, called “plasma relaxation”, and it can be understood in terms of the magneto-hydrodynamical (MHD) model of the plasma.
- The resulting magnetic field lines in the RFP follow helical paths along the torus. The magnetic field vector has components that are both poloidal (short way around the torus) (B_P) and toroidal (B_T) (long way around the torus). The plasma in the RFP tends to approach a state where the plasma electric current is everywhere aligned with the magnetic field, thus minimizing the magnetic Lorentz force. This results in a reversal of the toroidal magnetic field near the plasma boundary.
- The RFP plasma carries electric current. The toroidal current is driven inductively by pulsed external currents in a set of magnetic coils concentric with the plasma torus. The coil power supplies consist of multiple-stage capacitor banks, with banks discharged in sequence by timed closing switches.
- The poloidal plasma current arises through the relaxation process, in which the toroidal plasma current is “deformed” into a helical current path. In addition, there is a toroidal field coil maintaining the reversed magnetic field.




Project goals

- The project goal is to develop and carry out simulations of the electrical circuitry in EXTRAP T2R with the help of a simplified model of the RFP plasma relaxation.
- The RFP plasma is represented by a non-linear two-port circuit element, where the ports are connected to the external poloidal and toroidal field power supply circuits, respectively.
- Experimental waveforms of currents and magnetic fields from EXTRAP T2R will be used to find the circuit parameters for the plasma through comparison between experimental and simulated data.
- Suggested modelling tools are MATLAB and SIMULINK.

Context L: Solar wind and planetary environments

Context Responsible: Shane Carberry Mogan (ShaneCM@KTH.se)
Division 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 Hubble Space Telescope (HST) and the James Webb Space Telescope (JWST). The space plasma physics research group, SPP, 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 L1 uses magnetic field data from ESA's Cluster satellites to study how ultra-low frequency waves evolve into Short Large-Amplitude Magnetic Structures (SLAMS) at Earth's quasi-parallel bow shock.

Project L2 uses data from the NASA Cassini spacecraft to investigate variability in Titan's upper atmosphere, including signatures of waves, day-night and seasonal differences, and the effects of solar and magnetospheric conditions.

Project L3 analyzes data from NASA's Transiting Exoplanet Survey Satellite (TESS) to characterize host stars and then estimate atmospheric escape rates of orbiting planets, assessing whether they can retain potentially habitable atmospheres.

Project L4 uses new observations from JWST to search for and characterize atmospheric signals around Saturn's icy moons.

Project L5 develops code using NASA's SPICE toolkit to model solar eclipses of Jupiter's Galilean moons, calculating illumination patterns on their surfaces and generating visualizations to support telescope observations with Hubble and JWST.

Figure L6 analyzes data from spacecraft alignments in the inner heliosphere (e.g., Parker Solar Probe, Solar Orbiter, BepiColombo, Wind) to study how solar wind turbulence evolves with distance from the Sun and whether turbulent energy dissipation can account for solar wind heating.

Project L7 uses spacecraft data (e.g., Wind, Solar Orbiter) to test the scaling laws of cross-helicity in solar wind turbulence, measuring how it cascades and dissipates across scales and depends on solar wind conditions.

Project L8 sets up a Transient Array Radio Telescope (TART) on KTH campus, giving students hands-on experience with radio interferometry, hardware implementation, and preliminary sky observations using an open-source array.

Project L9 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 L1: The polarization of Short Large-Amplitude Magnetic Structures (SLAMS) at the bow shock of Earth

Supervisor: Sofia Bergman, sobergm@kth.se, Electromagnetics and Plasma Physics

Shock waves are found in many different environments throughout the universe. Some examples are shock waves of supernova remnants, planetary and stellar bow shocks and termination shocks around stars, pulsars and neutron stars. Most of these shocks can only be studied using remote methods, and therefore the shocks in our own solar system are frequently used as templates to understand other shocks in the universe.

The bow shocks of the planets are created when the supersonic solar wind (a stream of plasma emitted from the sun) interacts with the planets and is slowed to subsonic velocities. The nature of the bow shocks is highly dependent on the direction of the so-called interplanetary magnetic field (IMF). When the direction of the IMF is similar to the shock normal, the shock is referred to as quasi-parallel and is especially complex and turbulent. Quasi-parallel shocks have not been studied as much as other shocks due to their complex nature, and many questions remain regarding the physics taking place at these shocks.




Figure 1. Simulation of the bow shock of Earth [Karimabadi et al., 2014].

The region upstream of the quasi-parallel shock is filled with ultra-low frequency (ULF) electromagnetic waves. Also isolated non-linear structures, referred to as Short Large-Amplitude Magnetic Structures (SLAMS) are observed in this region (see Figure 2). SLAMS are believed to be very important for the formation of the quasi-parallel shock. They have been suggested to grow from the ULF waves, but the exact growth process is not known. One mystery is the polarization of these structures, which are observed to be opposite from that of the ULF waves.




Figure 2. SLAMS observed by Cluster.




Figure 3. The Cluster mission. Image from ESA.

The goal of this project is to study the polarization as ULF waves grow into SLAMS. You will use data from the Cluster satellites, which orbited the Earth between the years 2000 and 2024. Magnetic field data will be analysed to determine the polarization of SLAMS during different stages of their growth.

In this project, you will learn how to manage and analyse data from a space mission, including appropriate analysis techniques. You will also acquire general knowledge about the interaction of the solar wind with planets. The results will be very useful for future research by scientists at KTH.

Project L2: Exploring Titan's Upper Atmosphere with Cassini Spacecraft Data

Supervisor: Shane Carberry Mogan, ShaneCM@KTH.se, Electromagnetics and Plasma Physics

Titan, Saturn's largest moon, has a thick, nitrogen-rich atmosphere with surface pressures comparable to present-day Earth. Because Titan's mass and radius are much smaller, however, its atmosphere is far more extended, making it one of the most fascinating natural laboratories in the Solar System. Like the early Earth, Titan hosts complex organic chemistry and a methane-based weather cycle. Its upper atmosphere is constantly changing, influenced by internal processes such as waves, as well as by external forces, including solar photons and the charged particles of Saturn's magnetosphere.

Between 2004 and 2017, the NASA spacecraft *Cassini* carried out more than 100 close encounters of Titan. During these flybys, *Cassini's* Ion and Neutral Mass Spectrometer (INMS) directly measured gases in Titan's upper atmosphere. This produced a unique dataset spanning different chemical species, local times, solar cycle conditions, and seasonal changes over more than a decade. These data are still being analyzed today, providing an opportunity to study Titan's atmosphere in new and exciting ways.





Figure 1. (Left) Saturn and Titan as seen by the Cassini spacecraft. (Right) Molecular nitrogen (N_2) and methane (CH_4) measured in Titan's upper atmosphere by *Cassini*'s INMS between 2004 and 2015 (gray symbols), and inferred from *Cassini*'s Ultraviolet Imaging Spectrograph (UVIS) measurements during solar (black circles) and stellar (black triangles) occultations. From Müller-Wodarg & Koskinen (2025).

In this project, you will work with *Cassini*/INMS data to search for variability in Titan's upper atmosphere. The tasks include:

- Analyzing INMS data from several Titan flybys.
- Comparing measured density profiles to models to search for signatures of atmospheric gravity waves.
- Looking for diurnal differences (day vs. night, Titan's day lasts about 16 Earth days) and seasonal differences (each Saturnian season lasts about 7 Earth years).
- Testing whether the atmosphere looks different when Titan is exposed to the solar wind versus when it is shielded by Saturn's magnetosphere, and further exploring differences within the magnetosphere itself, such as when Titan is embedded in the plasma sheet compared to the magnetospheric lobes.
- Exploring whether changes in solar activity (the 11-year solar cycle) influenced *Cassini*'s measurements.

By carrying out this work, you will gain experience in spacecraft data analysis, statistical methods, and atmospheric physics, while contributing to ongoing planetary science research. The project also connects to future exploration: NASA's upcoming *Dragonfly* mission, which will land on Titan in the 2030s, will rely on improved understanding of Titan's atmosphere to guide its science goals. This project is well suited for students with interests in space physics, data analysis, and planetary exploration.

Project L3: Exploring Planetary Habitability with TESS Data

Supervisor: Shane Carberry Mogan, ShaneCM@KTH.se, Electromagnetics and Plasma Physics

The Transiting Exoplanet Survey Satellite (TESS) monitors large swaths of the sky, capturing high-precision light curves of stars. These data allow us to infer fundamental stellar properties such as luminosity, temperature, and flare activity. Once the radiation environment of the host star is known, we can investigate whether, and how quickly, planets orbiting those stars lose their atmospheres through atmospheric escape.

In this project, you will download and analyze TESS light curves from selected stars and carry out stellar characterization. You will estimate stellar luminosities and effective temperatures, and quantify flare activity (frequency and energy). With these properties in hand, you will calculate atmospheric escape rates for hypothetical planets orbiting the target stars. Using models of thermal (hydrodynamic) and energy-limited escape, you will explore how escape depends on orbital distance, planet mass and radius, and atmospheric composition. From this, you can determine whether a planet would retain its primordial hydrogen–helium envelope, and on what timescale. Planets that keep their envelopes evolve into worlds hostile to life, while those that lose them may develop secondary atmospheres, potentially more suitable for habitability.




Figure 1. (Left) Illustration of TESS observing an exoplanet system. Image credit: <https://www.cfa.harvard.edu/facilities-technology/telescopes-instruments/transiting-exoplanet-survey-satellite-tess>. (Right) Total energy released by flares from F-, G-, and K-type stars integrated over 100 million years corresponding to the flare's bolometric energy, E_{bol} , and frequency, ν_{flare} . Figure from Atri and Carberry Mogan (2021).

The scientific goal is to classify planets around the selected stars in terms of their habitability potential. This connects stellar astrophysics to exoplanet science, providing insight into how stars drive the long-term evolution of planetary atmospheres. This project offers hands-on experience in data analysis, stellar characterization, and atmospheric escape modeling. It is suitable for students interested in exoplanets, astrophysics, and computational modeling.

Project L4: James Webb Space Telescope observations of the icy moons of planet Saturn

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



Figure 1. James Webb Space Telescope (left) and Saturn's icy moons (right)

The NASA/ESA James Webb Space Telescope (JWST) was successfully launched on Christmas Day last year (2021). After commissioning of the telescope, science observations have been taken since July 2022 and the moons of Jupiter and Saturn were among the first targets for JWST.

JWST has the unique capabilities to take spectral images with information on wavelength (color) in infrared light (IR) in each pixel (Figure 2). Observations of Jupiter moons Ganymede and Callisto provided resolved images of the moons with information on the light's wavelength in each pixel. The data revealed exciting insights into the material in the moons' atmospheres and surfaces in different specific places. In this project, you will search for atmosphere signals in the JWST data of different moons of Saturn. Most of Saturn's moons are icy bodies and some might have oceans with liquid water below the icy crust. Atmospheres around the moons have been found by the NASA Cassini spacecraft but never seen in telescope observations before.

The tasks in this project include:

- Download JWST observations from one of the moons from the NASA data archive
- Read and process the data “cubes” extract infrared spectra and images of moons
- Interpret the spectra using reference spectra and simple models
- Compare the spectra to previous observations from spacecraft

Doing this project, you will be among the first people actively working with this milestone telescope in the world.




Figure 2. Example JWST observations showing emissions from CO₂ gas at Callisto.

Project L5: Solar eclipses seen from Jupiter's moons

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

Planet Jupiter has four large moons (Figure 1), with unique characteristics like extreme volcanic activity, huge subsurface water oceans, or an intrinsic magnetic field. Near equinox season at Jupiter, all four moons orbit through the shadow of Jupiter, i.e., they go through solar eclipse phases of several hours. This period is of particular interest to study the moons' surface and atmospheres, because the temperature drops very fast. The moons are regularly observed by telescopes during the eclipses including the ingress and egress phases, when they enter or leave the shadowed zone.



In this project, you will develop a code for the exact geometry of the solar eclipses experienced by Jupiter's moons. Specifically, the idea is to calculate the level of solar illumination at each point (planetary longitude and latitude) on the surface of a moon for a given time (from 100% for full Sun to 0% when the Sun is fully obscured by Jupiter). In addition, you can illustrate the appearance of the solar disk disappearing behind the disk of Jupiter as seen by an observer standing on a moon (a bit like shown Figure 1, middle, for the Earth's moon).

Steps include:


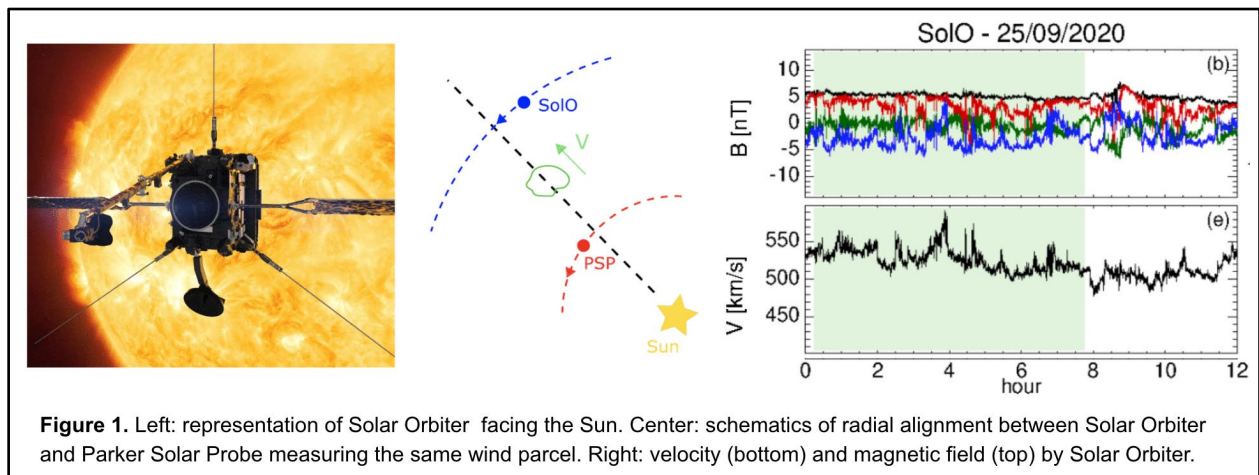
- Install and get familiar with the NASA SPICE software tool that allows you to calculate geometries for solar system objects (in MATLAB, Python, etc.)
- Produce a code that calculates to illumination as a function of coordinates and time for all four Galilean moons.
- Generate plots (similar to the one in Figure 1, right) that show the continuous level of illumination on the side of a moon that is observed by a telescope.
- Generate illustrations of the Sun as seen from a moon.

The resulting tool will be used to analyze observations from the NASA James Webb Space Telescope and Hubble Space Telescope. The students will be involved in the analysis of the telescope data, if this is of interest.

Project L6: Multi-spacecraft alignments and the radial evolution of solar wind turbulence

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

The inner heliosphere is currently sampled by several spacecraft, whose orbits occasionally result in the alignment of two or more of them along the radial direction from the Sun. Such configurations provide samples of expanding solar wind at different distances from the Sun, enabling us to evaluate the radial evolution of various properties. For this project, the student will use measurements collected during one or more radial alignments of spacecraft in the inner heliosphere, such as [Solar Orbiter](#), [Parker Solar Probe](#), [BepiColombo](#) or missions near the Earth ([Wind](#)). The accurate determination of the solar source region will be used to convalidate the effective radial alignment.



A comprehensive analysis of the turbulent statistical properties of the fields and plasma fluctuations will be complemented with estimates of the turbulent energy transfer rate, based on different versions of the third-order moment scaling laws.

The radial decay of the turbulent energy will be thus determined and compared with the measured solar wind heating, providing crucial information about the global energy budget of the solar wind in its expansion in the heliosphere.

In this project, the students will analyze data from spacecraft alignments in the inner heliosphere (e.g., Parker Solar Probe, Solar Orbiter, BepiColombo, Wind) to study how solar wind turbulence evolves with distance from the Sun and whether turbulent energy dissipation can account for solar wind heating.

Project L7: A cross-helicity cascade in solar wind turbulence: Statistical analysis of spacecraft data

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

The solar wind is a highly turbulent medium expanding from the sun to fill the whole heliosphere. As such, and thanks to spacecraft measurements, it serves as a solar-system-scale laboratory for fundamental studies of plasma turbulence. Due to its intrinsic unpredictability, turbulence theories often are based on a phenomenological description of statistical properties of the field fluctuations. There are only a few exact laws that prescribe exact relations, which are based on the conservation of invariants, such as the total energy or the magnetic helicity. One of such relations concerns the degree of alignment between velocity and magnetic fluctuations, quantified through a quantity called cross-helicity. While cross-helicity is dissipated by the turbulence in ways we do not yet understand, this project aims at validating the law describing the way it is degraded from the large-scale solar injection to the small-scale microscopic dissipation.



Figure 1. A computer simulation of the middle corona's magnetic field (Chitta 2024).

For this project, the students will use intervals of solar wind parameters measurements from spacecraft such as [Wind](#) and [Solar Orbiter](#), classify their turbulence properties, and compute the scaling laws that describe the cascade of cross-helicity.

The main questions that will be addressed are:

1. Is the cross-helicity scaling law valid in solar wind plasmas?
2. What is the cross-helicity dissipation rate in the solar wind?
3. How does the cross-helicity scaling and dissipation depend on the solar wind parameters?

In this project, the students will learn basic concepts of space plasma turbulence. They will acquire competences in obtaining, managing and analysing spacecraft data, and interpreting the observations of statistical analysis. The analysis will be performed using a programming language of choice, and possibly leveraging on previous work carried out by recent years' KTH students. The work will be part of a broad statistical study by KTH scientists and students.




Figure 2. Velocity and magnetic field fluctuations showing the alignment of solar wind conditions (Stumpo, 2025).

Project L8: Implementing a Transient Array Radio Telescope at KTH

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

Transient Array Radio Telescope ([TART](#)) is an international education/outreach project. The facility is a radiointerferometer array that observes the whole sky at once and produces an image by analyzing the phases of arrival of radiowaves. TART is an open-source project, allowing access and exploration of every part of a modern aperture synthesis radio telescope. This project aims to prepare a TART setup on KTH campus.

The tasks in this project include:

- Getting familiar with radio interferometric imaging.
- Implementing the hardware for the KTH instance of TART.
- Designing the setup for (semi-)permanent installation on campus.
- Making preliminary observations with TART.



Project L9: Observing pulsating aurora with a ground-based GNSS network

Supervisor: Judit Pérez-Coll Jiménez, juditpcj@kth.se, Nickolay 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 the magnetosphere, where trapped electrons interact with plasma waves and get diverted to the atmosphere, producing optical emissions. Apart from optical emissions, electrons 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.
- Analyzing 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.

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 Hubble Space Telescope (HST) and the James Webb Space Telescope (JWST). The space plasma physics research group, SPP, 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.

Context M: Machine Learning over Networks

Context Responsible: Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



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 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 **Machine Learning over Networks** comes into play. The challenge is to design 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 **Machine Learning over Networks**, 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.

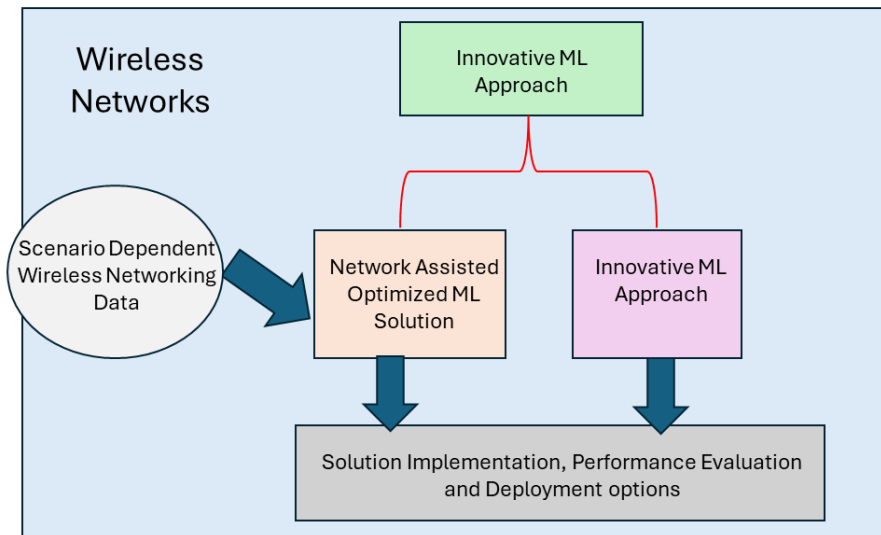
If you are passionate about AI, data science, and networking, this is your chance to be part of a transformative technological revolution through one of the following exciting projects:

- Over-the-Air Federated Learning
- An Introduction to Graph Neural Networks and Their Applications
- Predicting Network Latency Using Machine Learning
- Machine Learning Algorithms in Neural Networks
- Meta-Learning for AI-Aided IoT Networks
- Intelligent Over-the-Air Computation: Machine Learning Solutions for Edge Computing in IoT
- Distributed Optimization in Machine Learning
- Alternating Direction Method of Multipliers for Distributed Optimization
- Human Activity Recognition Using Machine Learning Through Wi-Fi Sensing

We describe the projects in detail on the following pages.

Project M1: Understanding Wireless Networks and Machine Learning (ML) to Develop Network-Assisted ML Solutions

Supervisors: Swapnil Sadashiv Shinde (ssshinde@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



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 M2: Understanding Key 6G Technologies and Their Integration for a Unified Framework

Supervisors: Swapnil Sadashiv Shinde (ssshinde@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering

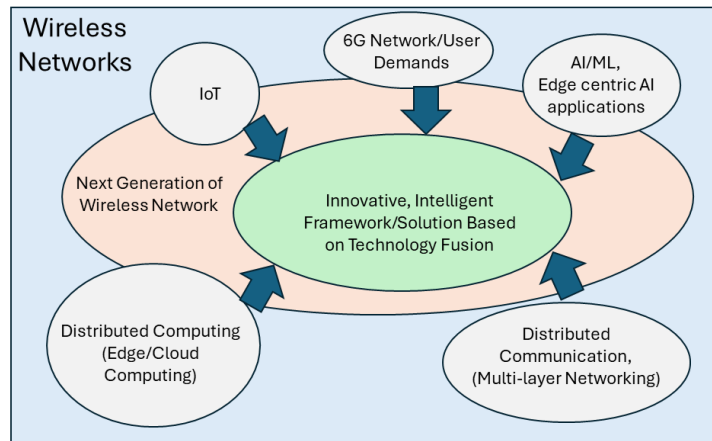


Figure 2: Unified 6G Framework Integrating Emerging Technologies

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**

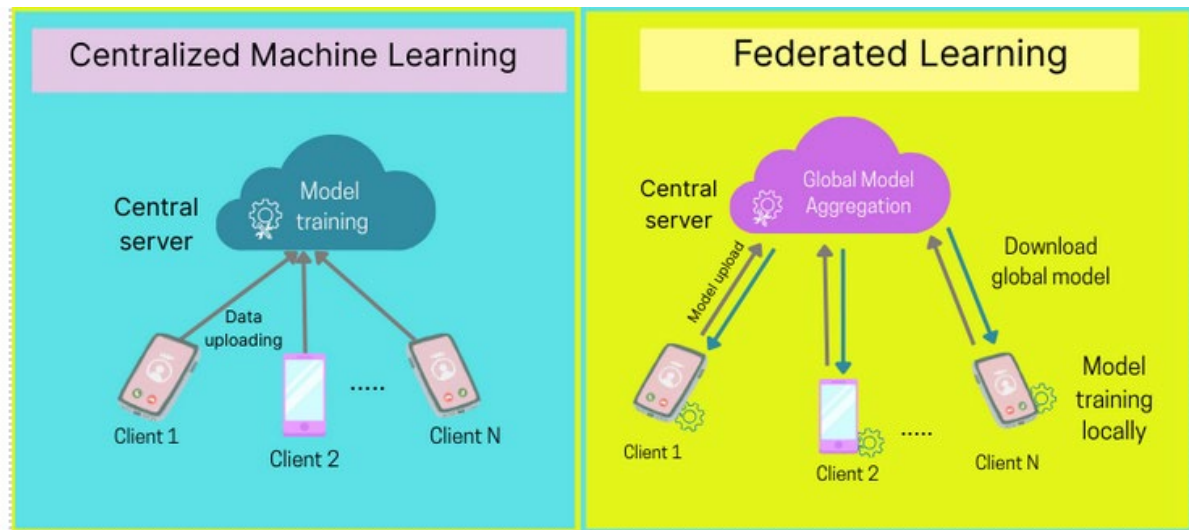
Understanding the requirements and interactions of these technologies is essential for designing a **unified 6G framework** that ensures their efficient coexistence and integration. In this project, the student will focus on **learning, analyzing, and summarizing existing knowledge** about 6G technologies 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 M3: Distributed Image Classification with Federated Learning

Supervisors: Milica Jankov (milicaj@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering




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 M4: First-Order Optimization Methods for Machine Learning

Supervisors: Seyedamirreza Kazemi (seykaz@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



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 M5: Online Convex Optimization

Supervisors: Krishnendu S. Tharakan (tharakan@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



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.

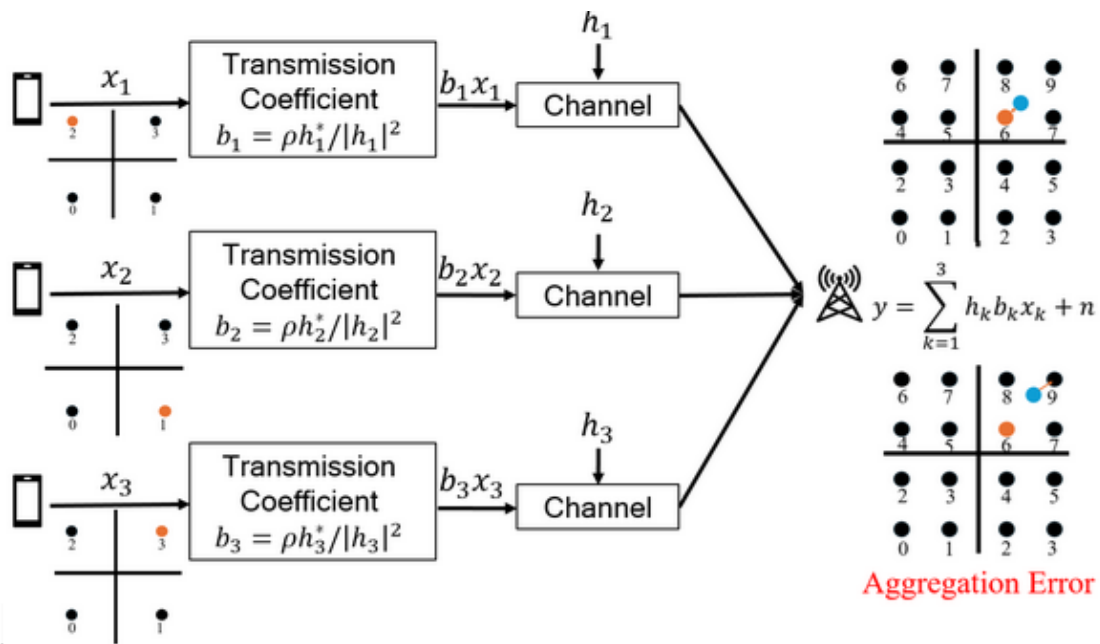
This project explores the online learning version of the knapsack problem with switching costs. This project explores fundamental algorithms in OCO, such as Online Gradient Descent (OGD) and Follow-The-Regularized-Leader (FTRL), and evaluates their theoretical guarantees and practical performance.

- 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 M6: Digital Modulation for over the air computations

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

Division of Network and Systems Engineering



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 M7: Semantic-Aware GNNs for Networked Vehicles

Supervisors: Minyi Wan (minyiw@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



Graphstructured 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, vehicularflavored graph for semantic communication.

The main objectives are as follows:

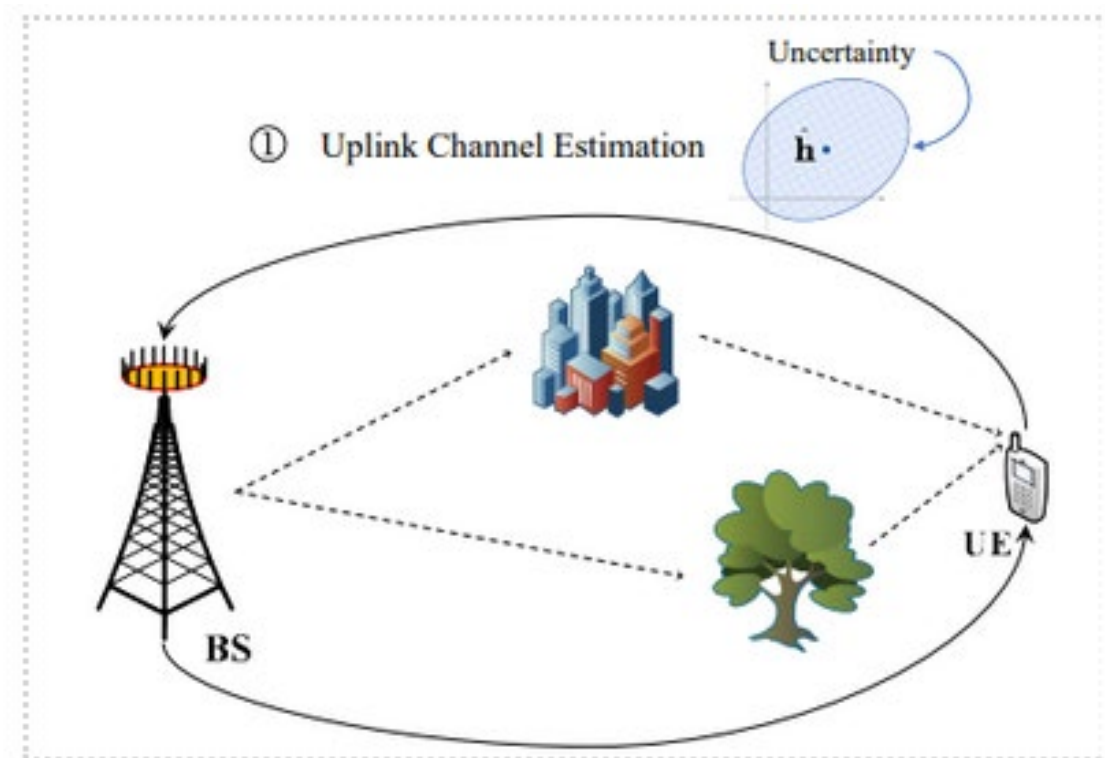
- Literature Review:
- Conduct a literature review of GNNs algorithms, including basic GNN, GCN, GAT, etc.
- Study general ideas about semantic communication concepts, understand the difference with three levels from Shannon-Weaver Explanation [1].
- Model Construction
- Build and simulate a simple GNN (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.

Reference:

- [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 M8: Conformal Prediction for Reliable Channel Estimation

Supervisors: Xin Su (xisu@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



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 M9: Distributed Optimization in Machine Learning

Supervisors: Seif Hussein (seifh@kth.se), Carlo Fischione (carlofi@kth.se)
Division of Network and Systems Engineering



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

Context N: Cyber Security

Context Responsible: Hamed Nemati (hnnemati@kth.se)
Division of Network and Systems Engineering



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. Intrusion detection and intrusion prevention
2. Active learning for intrusion detection and response
3. Trustworthy Autonomy in Cyber-physical Systems

Project N1: Automated Intrusion Response for IT systems

Supervisor: Rolf Stadler (stadler@kth.se), Duc Huy Le (dhle@kth.se), Division of Network and Systems Engineering



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 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 N2: Designing and Simulating Cyber Attacks on Drones

Supervisors: Axel Andersson (axander@kth.se), György Dán (gyuri@kth.se), Division of Network and Systems Engineering



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 by 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 N3: Adversarial Attacks on Multi-Agent Reinforcement Learning Systems

Supervisors: Kiarash Kazari (kkazari@kth.se), György Dán (gyuri@kth.se), Division of Network and Systems Engineering



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 N4: Robustness of RNN-Based World Models

Supervisors: Manjodh Singh (manjodh@kth.se), György Dán (gyuri@kth.se), Division of Network and Systems Engineering



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.

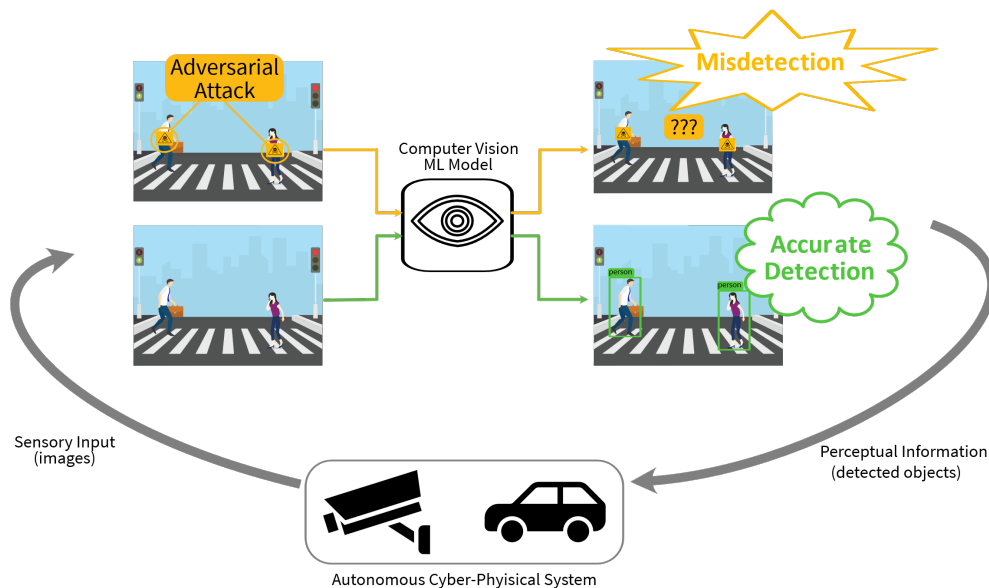
For more background information about the project, (Ha and Schmidhuber, 2018) and (Zeng et al. 2024) serve as good starting points for exploring the rich literature on world models.

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 N5: Trustworthy Autonomy in Cyber-Physical Systems

Supervisors: Mauricio Byrd Victorica (mbv@kth.se), György Dán (gyuri@kth.se), Division of Network and Systems Engineering



The performance and reliability of autonomous Cyber-Physical Systems (CPS) depend on (i) the accuracy of both the sensory information they acquire (e.g., camera) and (ii) the AI/ML algorithms they use to process it (e.g., object detection models). Adversarial attacks are sensory perturbations designed to mislead ML models, and thus pose a safety risk for autonomous CPS [1, 2, 3].

Early works mainly focused on attacks and defenses for ML-models in isolation, but recent works have shifted towards adversarial attacks which are both physically-realizable and capable of decision-level impact on full-fledged systems where the ML-model under attack is only a component of the system's perception pipeline (e.g., autonomous vehicles) [4]. In turn, defense mechanisms to handle attacks with system-level impact have also been proposed [5].

The objective of this project is to improve the trustworthiness of CPS that rely on ML for perception (e.g., ML computer vision and vision-language models). This might involve, among other tasks:

- Creating adversarial perturbations, i.e., attacks.
- Designing mechanisms to defend ML models from adversarial attacks.
- Modifying and/or integrating datasets or complex simulation environments to evaluate attack and defense methods in realistic scenarios [attack on autonomous driving demo].

Related Work:


[1] Eykholt et al, "Physical Adversarial Examples for Object Detectors," in Proc. of WEET, 2018.

[2] Mauricio Byrd Victorica et al., "SpaNN: Detecting Multiple Adversarial Patches on CNNs by Spanning Saliency Thresholds" in Proc. of SaTML, 2025.

[3] Mauricio Byrd Victorica et al., "SaliuIt!: Ensemble Saliency Guided Recovery of Adversarial Patches against CNNs", in Proc. of CVPR, 2025.

[4] Yi Zhu et al. "Malicious Attacks against Multi-Sensor Fusion in Autonomous Driving", in Proc. of ACM MobiCom, 2024.

[5] Raymond Muller et al. "VOGUES: Validation of Object Guise using Estimated Components", in Proc. of USENIX Security, 2024.



Context Responsible: Tobias Oechtering (oechtering@kthlse)
Division of Information Science and Engineering

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.

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 department's teachers. Accordingly, all offered Bachelor projects are closely related to some of the on-going research projects in the division. In particular, we offer project that deal with data augmentation for ECG training data, algorithms for contactless heart-beat sensing, assessment of data pre-processing methods, tuning of GPT model of a chess bot, machine learning based FM radio receiver, and model inversion attack-based assessment of privacy-enhancing subsampling. Since information processing is quite abstract, all projects require a good mathematical background and solid programming skills.

Project O1: Synthetic Heart Sound Using Generative Models

Supervisor: Xinqi Bao, xba@kth.se, and Saikat Chatterjee, sach@kth.se, Division of Information Science and Engineering

Problem Statement

Heart sounds (phonocardiograms, PCG) are an important non-invasive indicator of cardiac structure and function. Automatic recognition of heart sounds can assist in early screening for structural heart diseases. However, publicly available PCG datasets are small, class distributions are highly imbalanced, and annotation is costly because physician auscultation and confirmation are required.

Synthetic heart sound generation using modern generative models (such as GANs and Diffusion models) can address these limitations by:

- Enabling data augmentation for rare pathological classes.
- Supporting medical education and training with realistic heart sound samples.
- Protecting patient privacy by generating non-identifiable data.

Although generative methods are widely explored for ECG and audio, their application to PCG is still in its infancy. This thesis will investigate whether GAN and Diffusion models can generate realistic heart sounds and how these synthetic signals can benefit classification models



Figure 1. A heart sound auscultation

Aim of the Thesis Project

The aim is to design, implement, and evaluate **generative models for synthetic heart sound** using open datasets (PhysioNet Challenge 2016 and 2022). The project will:

1. Develop and compare at least one generative model for PCG synthesis.
2. Evaluate the realism of generated signals using quantitative metrics (e.g., Mel-Cepstral Distortion, Pearson correlation, KL divergence) and qualitative assessment (blind listening or visual inspection).
3. Test whether the synthetic data improves a baseline heart sound classification model.

Optional advanced goals include multi-channel/multi-site heart sound generation and conditional generation (e.g., by murmur type or heart rate).

Required Skills

- Good programming skills in **Python** (PyTorch or TensorFlow).
- Familiarity with **signal processing** and **time-series/audio data**.
- Understanding of **machine learning concepts**; prior exposure to GANs or Diffusion models is a plus but not mandatory.
- Interest in biomedical signal analysis and generative models.

References

- Clifford GD, et al. (2016) **PhysioNet Challenge 2016: Classification of Heart Sound Recordings**. In *2016 Computing in cardiology conference (CinC)* (pp. 609-612). IEEE.
- Reyna, MA, et al. (2022) **Heart murmur detection from phonocardiogram recordings: The george b. moody physionet challenge 2022**. In *2022 Computing in Cardiology (CinC)* (Vol. 498, pp. 1-4). IEEE.
- Alm Narváez, P. and Percybrooks, WS, (2020). **Synthesis of normal heart sounds using generative adversarial networks and empirical wavelet transform**. *Applied Sciences*, 10(19), p.7003.
- Wang, Z., et al., (2024) **ECG Generation Based on Denoising Diffusion Probabilistic Models**.

Project O2: Straggler-Resilient Distributed Learning and Computer Vision for Medical Diagnosis

Supervisor: Jeannie He, jeannie@kth.se, and Ming Xiao, mingx@kth.se, Division of Information Science and Engineering

Problem Statement:

Recent advances in research have highlighted the need for distributed learning algorithms that can handle the challenges of large-scale, real-world deployments. For example, in [1], a straggler-resilient asynchronous ADMM algorithm was proposed to enable distributed learning whilst mitigating delays from slower (straggling) nodes. In this project, the aim is to integrate this asynchronous ADMM method with a state-of-the-art image processing model, such as YOLOv8, a state-of-the-art image analysis model. The specific application is medical diagnosis, where distributed learning offers several advantages:

1. Privacy preservation – Medical data is highly sensitive, and hospitals or clinics are often restricted from sharing raw patient images due to legal and ethical concerns. Distributed learning enables institutions to collaborate without moving patient data outside their own secure environments.
2. Improved reliability – In medical settings, not all participating institutions will have equal computational resources. Some may be “stragglers” that slow down training in standard federated learning setups. An asynchronous, straggler-resilient approach allows the system to continue learning effectively despite these variations.
3. Enhanced diagnostic accuracy – By leveraging YOLOv8’s advanced object detection and segmentation capabilities, the system can learn from diverse medical images (e.g., X-rays, CT scans) collected across multiple hospitals, leading to models that generalize better across patient populations.
4. Scalability for real-world healthcare networks – A method that combines robustness, privacy, and efficiency can make distributed AI solutions practical for large medical collaborations, where performance and trustworthiness are critical.

Specifically, the student will:

1. Implement and incorporate YOLOv8 into the straggler-resilient asynchronous ADMM algorithm [1] based on available source codes,
2. Evaluate the performance (accuracy, convergence speed, robustness to stragglers, communication overhead) on open source MRI images,
3. Compare the performance of this hybrid framework against YOLOv8 with conventional federated learning (FL) and/or other frameworks.

You may use open source code for the project. You will get access to the source code.

Required skills and knowledge: Python, Machine learning

References:

[1] He, J., Xiao, M., Skoglund, M., & Poor, H. V. (2025). Straggler-Resilient Asynchronous ADMM for Distributed Consensus Optimization. *IEEE Transactions on Signal Processing*.

Project O3: Explainable Machine Learning for Health Care

Supervisor: Ragnar Thobaben (ragnart@kth.se), Division of Information Science and Engineering (ISE)

In the last decade, a large amount of health care data has been collected in hospitals worldwide, which is now partially made available for research purposes in public data bases to foster the development of new data-driven approaches to health care and to take advantage of the recent progress in the field of machine learning. Typical applications for machine learning in health care include pre-processing steps (e.g., outlier detection, missing-data imputation), decision support systems, (e.g., diagnostics, outcome prediction, and treatment recommendations), treatment automation (e.g., glucose, ventilation, circulatory and cardiovascular management), and data analysis (e.g., identification of patient- and treatment-specific risk factors). One of the biggest obstacles though that hinders the wide acceptance of machine learning models in health care, is lack of interpretability of such models. For example, in a clinical context, a decision support system needs to be able to provide an explanation for its recommendation in order to be trusted by the medical staff, and in the case of data analysis (e.g., identification of risk factors), it is important to understand which signal components are responsible for triggering a certain response of the learned model. In this project, we will investigate this issue by studying machine learning models trained from health-care data (e.g., models for detecting cardiovascular disease and/or pre-diabetes, mortality prediction) and by extending these models to yield interpretability and explainability. This can be achieved, e.g., by considering models that are inherently well interpretable like, e.g., logistic regression, random forests, explainable boosting machines, and clustering, or by performing a sensitivity analysis of trained models with known parameters.

The steps in the project are as follows:

- In self-studies, the students are expected to acquire the required machine learning background and skills to execute the work in this project. Following publicly available online lectures has been a successful approach to this in previous years.
- Together with the supervisor, you will identify a few suitable problems (e.g., detection of cardiovascular disease and/or pre-diabetes, mortality prediction) and data sets, for which you will train at least two different machine learning models using different methods.

- Next, you will adopt strategies from the literature and possibly develop new strategies to further analyse the models to explain their responses to the data.
- You will summarize your findings and present the results of a comparison of the different approaches investigated in this project in the final report and the final presentation. The final report will also include a brief survey of recent approaches to explainable machine learning.

The project is fairly open and leaves students with a lot of space to develop and pursue own ideas. Since this freedom also can be a burden, this project is only recommended for creative students with strong mathematical and programming background. Students working in this project will be supported by Ragnar Thobaben.

Project O4: Hierarchical Inference Learning for Object Boundary Detection via Image Segmentation

Supervisor: Adarsh Behera, apbehera@kth.se, and James Gross, jamesgr@kth.se,
Division of Information Science and Engineering

Introduction: Hierarchical Inference (HI) is a meta learning framework that uses reinforcement learning based expert-advice prediction to balance the use of lightweight, embedded machine learning models on edge devices (SML) with more powerful models at edge servers (LML). See Fig.1. By learning when to accept local predictions and when to offload for higher accuracy, HI provides a trade-off between accuracy, latency, energy efficiency, and bandwidth usage. Recent algorithms extend HI using online learning methods, enabling resource-constrained devices to adaptively decide offloading thresholds without prior knowledge of data distributions. While the concept can be generalized, these are applied, until now, on image classification and event detection usecases.

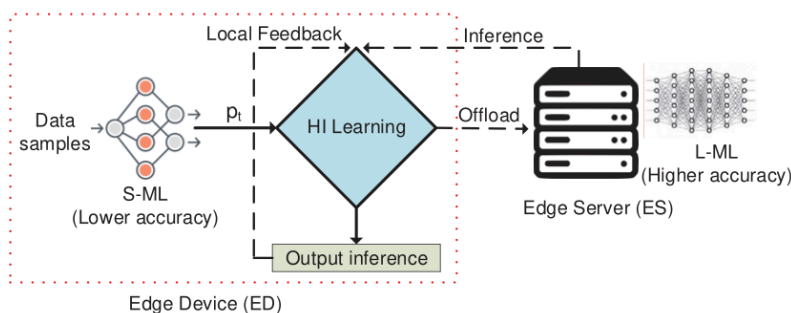


Figure 1: Basic Hierarchical Inference

Problem Statement: In this project, the aim is to extend the HI framework to the problem of **image segmentation for object boundary detection**. Unlike simple classification, segmentation requires precise boundary localization, where confidence is based on how accurate the boundary that is detected. Thus, an important task here will be to design/study/adapt a **confidence metric** for segmentation predictions (e.g., boundary uncertainty, or region overlap measures), and then integrate this metric into the HI decision process. The modified HI algorithm should learn online whether to accept a local segmentation boundary or to offload for refinement. A starting point would be to consider only a single contiguous boundary scenario, and use the simplest of HI algorithms, referred to as HIL-F in literature, that assumes the availability of ground truth at SML after the decision is made.

Aim of the thesis project: The project aims to implement and evaluate an extended HI algorithm for image segmentation. Students will:

- Propose and implement a confidence metric suitable for boundary detection tasks.
- Modify the HIL-F to incorporate this metric into the offloading decision.
- Find suitable datasets, find the confidence metric Evaluate the algorithm on such datasets comparing trade-offs in accuracy, offload rate, and resource efficiency.
- If time permits, extend the results to non-contiguous segmentation and HIL-N.

Note that prebuilt and readymade datasets and segmentation models are okay. This project is not about segmentation but about using RL-based HI approach to learn which of the two segmentation models needs to be used at inference time.

Primary basic software skills required: MATLAB and/or python.

References:

- [1] V. N. Moothedath, J. P. Champati and J. Gross, "Getting the Best Out of Both Worlds: Algorithms for Hierarchical Inference at the Edge," in *IEEE Transactions on Machine Learning in Communications and Networking*, vol. 2, pp. 280-297, 2024, doi: 10.1109/TMLCN.2024.3366501.
- [2] Letsiou, Afroditi, et al. "Hierarchical inference at the edge: A batch processing approach." *2024 IEEE/ACM Symposium on Edge Computing (SEC)*. IEEE, 2024.

Project O5: Hierarchical Inference Learning for Event Detection in Time Series

Supervisor: Afroditi Letsiou, aletsiou@kth.se, and James Gross, jamesgr@kth.se, Division of Information Science and Engineering

Introduction: Hierarchical Inference (HI) is a meta learning framework that uses reinforcement learning based expert-advice prediction to balance the use of lightweight, embedded machine learning models on edge devices (SML) with more powerful models at edge servers (LML). See Fig.1. By learning when to accept local predictions and when to offload for higher accuracy, HI provides a trade-off between accuracy, latency, energy efficiency, and bandwidth usage. Recent algorithms extend HI using online learning methods, enabling resource-constrained devices to adaptively decide offloading thresholds without prior knowledge of data distributions. While the concept can be generalized, these are applied, until now, mostly on image classification use cases.

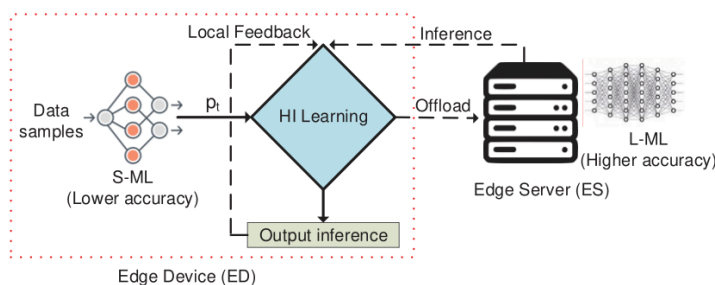


Figure 1: Basic Hierarchical Inference

Problem Statement: In this project, the aim is to extend the HI framework to the problem of event detection through monitoring a time series dataset. Unlike static classification tasks, event detection in time series requires identifying the occurrence of events within a temporally correlated stream, where confidence depends not only on the current prediction but also on its consistency over time. This motivates the need for designing or adapting confidence measures tailored to temporal signals (e.g., stability across sliding windows, or prediction variance), which can then be embedded into the HI framework to decide when local detections are

sufficient and when offloading is necessary. The study can focus on a simple scenario with one type of event, and explore how the online HI algorithms, such as HIL-F or HIL-N, can be modified to handle the characteristics of event detection in time series.

Aim of the thesis project: The project aims to implement and evaluate an extended HI algorithm for time series. Students will:

- Study and identify suitable confidence metrics for event detection in time series.
- Extend the existing hierarchical inference algorithms to incorporate the need of the different confidence metric
- Evaluate the adapted framework on one or two real world time series datasets by measuring accuracy, detection delay, and communication overhead.

Note that, the main objective of this project is not to improve the event detection method itself, but rather to investigate what modifications are required in the online HI algorithms so that they can be tailored to different types of ML tasks, with a particular focus on event detection in time series.

Primary basic software skills required: MATLAB and/or Python.

References:

[1] V. N. Moothedath, J. P. Champati and J. Gross, "Getting the Best Out of Both Worlds: Algorithms for Hierarchical Inference at the Edge," in *IEEE Transactions on Machine Learning in Communications and Networking*, vol. 2, pp. 280-297, 2024, doi: 10.1109/TMLCN.2024.3366501.

[2] Letsiou, Afroditi, et al. "Hierarchical inference at the edge: A batch processing approach." *2024 IEEE/ACM Symposium on Edge Computing (SEC)*. IEEE, 2024.

[3] Kharazian, Zahra, et al. "SCANIA component X dataset: a real-world multivariate time series dataset for predictive maintenance." *Scientific Data* 12.1 (2025): 493.

[4] Scania, C. V. A. B. "Air pressure system failures in Scania trucks." URL: <https://www.kaggle.com/uciml/aps-failure-at-scania-trucks-data-set>. (accessed 11.30. 2021)

Project O6: Signal Integrity for Early NICU Alerts: Artifact Detection, Missingness Modeling, and Robust Preprocessing

Supervisor: Zhendong Wang, Yogesh Todarwal and Saikat Chatterjee, sach@kth.se, Division of Information Science and Engineering

Background and Significance

This is a health-care and medical data analysis project in collaboration with Karolinska Institute. Early-warning models for neonatal adverse events depend critically on the integrity of raw bedside signals. Motion artifacts, sensor clipping, clock drift, irregular sampling, and missingness can degrade signal quality, bias features, and inflate false alarms. This project focuses on the data-centric pipeline prior to modeling: rigorous signal quality assessment, artifact detection and repair, missingness analysis and imputation, resampling/synchronization, and denoising. The goal is to develop a preprocessing framework that measurably improves downstream early-warning alerts.

Objectives and Hypotheses

Objective:

Design and evaluate a reproducible preprocessing pipeline for NICU vital signs that improves signal integrity prior to modeling.

Hypotheses:

1. Automated artifact detection and repair will reduce spurious volatility in derived features and lower false-alarm rates in a fixed downstream model.
2. Missingness-aware representations (mask vectors and time-since-last-observed features) combined with principled imputation will increase the effective sample size without biasing event timing.
3. Proper resampling, synchronization, and denoising will increase feature stability and yield earlier and more reliable early-warning predictions.

Methodology**1) Artifact Detection & Repair.**

Design rule-based and statistical detectors for sensor clipping/plateaus, abrupt jumps/drops, flatlines, out-of-physiology ranges, and step changes. Implement repairs via masking, local interpolation, and median/Hampel filtering. Provide per-window quality scores.

2) Missing Data

Create a binary mask and a time-since-last-observed feature for each signal. Short gaps (≤ 10 – 15 s) will be filled with linear or cubic interpolation; medium gaps (10 – 120 s) with Kalman smoothing; long gaps (>120 s) will remain missing. Data near diagnosis times will not be imputed. The optimal method is chosen by a hide-and-fill test on clean segments (hide known points, fill them, and select the lowest-error method). All imputations are logged (fraction filled, gap lengths, method).

3) Resampling, Synchronization, and Drift Correction

Unify sampling grids across channels; assess aliasing and lag effects. Implement clock-drift checks via cross-correlation and align per-patient timelines.

4) Denoising & Robust Feature Extraction

Evaluate various denoising pipelines (e.g., moving median, Hampel). Assess the stability of clinically important features (e.g., entropy) under noise and artifacts. Report a stability summary (variance reduction, outlier rate).

5) Evaluation Protocol

- Intrinsic (signal-level): Percentage of data retained, number of artifacts removed, SNR improvement, error in hide-and-fill tests, and residual drift after synchronization.
- Downstream (model-level): Apply a previously trained classifier (without retraining) to data before and after preprocessing; compare AUROC/AUPRC and lead time at fixed specificity. Improvements can be attributed to preprocessing alone.

Ethical & Data Governance Considerations

Use de-identified data under existing approvals and keep data processing within secure environment.

Required Skills & Learning Outcomes

Skills: Python, time-series signal processing, imputation fundamentals, scientific writing. Outcomes: Mastery of data-centric pipelines for physiological time series; rigorous evaluation of signal integrity; reproducible reporting.

Context P: 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 P1: 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 AAI 2012 (2012)


Project P2: Anti-Submarine Warfare with Counterfactual Regret Minimization

Handledare: 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), 2025.

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 P3: Beyond Visual Range Air Combat with Monte-Carlo Tree Search

Handledare: 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, 2025

The Animated Monte-Carlo Tree Search (MCTS), Kurbiel, <https://medium.com/data-science/the-animated-monte-carlo-tree-search-mcts-c05bb48b018c>.

Project P4: Explainable Fuzzy Challenge 2026

Handledare: 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 avfyra 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 2025 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 2026.

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/I3kN0NRlYYA>
- Free Optimization, <https://youtu.be/6qfPmJHaBNE>

Context Q: 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 Q1: 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 Q2: Brain-like neural networks and associated neuromorphic computing algorithms – 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 R: 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 R1: Sign Spotting

Supervisor: Jonas Beskow (beskow@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 R2: Recognizing signing attributes using vision-language models

Supervisor: Jonas Beskow (beskow@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 R3: Sign language detection

Supervisor: Jonas Beskow (beskow@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. from 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 R4: Masked sign language prediction model

Supervisor: Jonas Beskow (beskow@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.

Project R5: Exploring rotation parameterization effects on sign language production quality

Supervisor: Anya Klezovich (annkle@kth.se)

Sign Language Production (SLP) refers to generation of sign language from a semantic representation (aka synthesis). In this project you will set up the standard SLP baseline, based on progressive transformers SLP paper [1], and then compare three rotational pose encodings for SLP: unit quaternions, exponential maps, and the continuous 6D rotation representation [2], extending on the original paper Fauré et al. [3]. Fauré et al. [3] reported improved performance from quaternion-based geodesic losses and contrastive supervision but did not compare exponential maps or 6D rotations. As a result of this project it will become more clear which rotation parameterization best balances positional accuracy and angular fidelity for SLP when working with the 3D data.

This project requires previous experience of training models.

[1] Saunders, B., Camgoz, N. C., & Bowden, R. (2020, August). Progressive transformers for end-to-end sign language production. In European Conference on Computer Vision (pp. 687-705). Cham: Springer International Publishing.

[2] Zhou, Y., Barnes, C., Lu, J., Yang, J., & Li, H. (2019). On the continuity of rotation representations in neural networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 5745-5753).

[3] Fauré, G., Sadeghi, M., Bigeard, S., Ouni, S. (2025). Towards Skeletal and Signer Noise Reduction in Sign Language Production via Quaternion-Based Pose Encoding and Contrastive Learning. In ACM International Conference on Intelligent Virtual Agents (IVA Adjunct '25), September 16–19, 2025, Berlin, Germany. ACM, New York, NY, USA, 9 pages.

<https://doi.org/10.1145/3742886.3756728>