Aalto Science Institute (AScI)

International Summer Research Programme

2026 project list

(updated, 07.01.2026)

For more information on the program and how to apply, see
https://www.aalto.fi/en/aalto-science-institute-asci/aalto-science-institute-international-summer-research-programme

Note that many projects have multiple positions.

Contents

SCHOOL OF CHEMICAL ENGINEERING 3

Department of Chemistry and Materials Science **3**
1201 - Sustainable materials and devices for green transition 3

SCHOOL OF ELECTRICAL ENGINEERING 4

Department of Electrical Engineering and Automation **4**
2101 - Learning and Control for Networked Multi-agent Systems 4
2102 - Controllable Medical Synthetic Data with Constrained Diffusion Models 5
2103 - Development of Advantage-Based Reinforcement Learning Algorithms 6

Department of Electronics and Nanoengineering **7**
2201 - Applications of quantum computing in optimization 7

Department of Information and Communications Engineering **8**
2301 - Design and Implementation of an Administrative Web Front-End for AI Summarization 8
2302 - Simulating how people read with reinforcement learning 9
2303 - AI alignment by combining different forms of Human Feedback 10
2304 - Uncertainty-Aware Meta-Learning for User Modelling 11
2305 - Human-in-the-Loop Bayesian Optimization for Design 12
2306 - AI of the Beholder: VLM-based Evaluation of Character Movement Quality 13
2307 - Bayesian Optimization for Eliciting Individuals' Creative Potential 14
2308 - Simulating Health Behaviors to Design Interventions Through Reinforcement Learning 15
2309 - UniHand: Unified Hand Control 16

SCHOOL OF ENGINEERING 17

Department of Civil Engineering **17**
3201 - EcoNORDICS: A focus on developing Science-based, Eco-Construction Practices 17
3202 - Digital Construction- Computer Vision perspective 18
3203 - Simulation of Construction Processes – Digital Twin Perspective 19

Department of Applied Physics **20**
4101 - Biofabricated mycelium cellulose cellular solids for programmable insulation and structures 20
4102 - Photon transport in fusion plasma boundaries 21
4103 - Tensor-Network Tight-binding Hamiltonian for fractals 22
4104 - Radiation damage modelled with two-temperature molecular dynamics 23
4105 - Dark optical-Paul trapping of large mass particles in ultra-high vacuum / Entangled light-matter states in polariton optomechanical systems 24

Department of Computer Science **25**
4201 - A Principled Approach to Watermarking Diffusion Language Models 25
4202 - Computational Perspectives on Democracy 26
4203 - Machine Learning: Foundations and New Frontiers 27
4204 - Social Intrinsic Motivation 28
4205 - Bayesian physician — training a Bayesian neural network to emulate physician's decisions 29
4206 - Geometric algorithms in Euclidean and hyperbolic spaces 30
4207 - Probabilistic machine learning and generative AI 31
4208 - Distributed quantum computing 32
4209 - Massively Parallel Algorithms for Graph Problems 33
4210 - Bayesian workflow 34
4211 - The psychology of Human-AI interaction 35
4212 - Designing New Intimate Wearables for Reproductive Health 36
4213 - Analysing Network Support for Women in Scientific Careers 37
4214 - Change in web browsing behavior across shock events 38

School of Chemical Engineering

Department of Chemistry and Materials Science

1201 - Sustainable materials and devices for green transition		
Field of study:	Chemistry	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	2	
School:	School of Chemical Engineering	
Department:	Department of Chemistry and Materials Science	
Professor:	Tanja Kallio	tanja.kallio@aalto.fi
Academic contact person:	Milla Vikberg	milla.vikberg@aalto.fi

Are you interested in developing more sustainable materials and devices for such applications as lithium-ion batteries or hydrogen generation via water electrolysis? Our team focuses on improving material and energy efficiency of these and other emerging electrochemical energy conversion and storage processes to facilitate the ongoing green transition.

The practical research work covers material synthesis, chemical characterizations, and electrochemical investigations in lab-scale devices and aims at understanding the connection between material structure and electrochemical properties. If you are interested in the positions, previous experience in synthesis and characterization is beneficial, likewise knowledge on electrochemistry.

School of Electrical Engineering

Department of Electrical Engineering and Automation

2101 - Learning and Control for Networked Multi-agent Systems

Field of study:	Automatic control, machine learning	
For students currently studying:	Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering	
Department:	Department of Electrical Engineering and Automation	
Professor:	Dominik Baumann	dominik.baumann@aalto.fi
Academic contact person:	Dominik Baumann	dominik.baumann@aalto.fi

In the next generation of control systems, we will see multiple smart agents that collaborate to reach a joint goal. For instance, autonomous cars will exchange their planned routes to increase traffic throughput, drone swarms will coordinate to increase efficiency in precision agriculture, and mobile robots will collaborate with factory automation machinery to manufacture products jointly. To achieve this level of collaboration, the agents need to be able to (i) autonomously act in and learn from interactions with their environment, and (ii) communicate with each other via a wireless communication network. Thus, the next generation of control systems will integrate classical control theory with machine learning and wireless networks. While both learning and wireless communication offer unprecedented flexibility, they at the same time introduce novel challenges. Wireless networks have a limited bandwidth, and abundant use of the network through communicating unnecessary information will lead to delay and even loss of sent messages. Learning requires the system to explore and try out actions. However, for application examples like mobile robots or autonomous cars that act around humans, strict safety guarantees are indispensable. Providing such guarantees has always been a major strength of control theory. Nevertheless, existing techniques must be enhanced to include the communication and learning system.

In this project, we seek to develop algorithms that allow learning and control of multi-agent systems over communication networks while providing safety guarantees. The concrete tasks within the project can be adjusted based on the background and interest of the applicant. The focus can be on the mathematical analysis of the developed algorithms, an evaluation in complex simulation environments, or an evaluation on a real robotic system, e.g., a Franka robot arm or the Unitree Go2 quadruped.

2102 - Controllable Medical Synthetic Data with Constrained Diffusion Models

Field of study:	Generative AI; Machine Learning; Medical Imaging; AI for Healthcare	
For students currently studying:	Master's or PhD	
Number of positions offered:	1-2	
School:	School of Electrical Engineering	
Department:	Department of Electrical Engineering and Automation	
Professor:	Jiancheng Yang	jiancheng.yang@aalto.fi
Academic contact person:	Jiancheng Yang	jiancheng.yang@aalto.fi

High-quality synthetic medical data is becoming increasingly important for modern AI systems due to privacy restrictions, limited annotations, domain shift, and access barriers in clinical environments. Recent diffusion models have shown impressive generative capabilities, yet achieving clinically meaningful, anatomically consistent, and controllable medical data generation remains an open research challenge.

This project investigates constrained diffusion models for medical synthetic data generation, where medical priors, anatomical structures, or task-specific constraints are explicitly embedded into the generative process. The objective is to develop diffusion-based models capable of producing high-fidelity, controllable, and clinically valid synthetic medical images or multimodal datasets. Applications include data augmentation, privacy-preserving model training, improved robustness, and downstream diagnostic tasks.

The summer intern(s) will have the opportunity to:

- Implement and train diffusion-based generative models, including integrating medical or physical constraints
- Work with ample compute resources (ELLIS Institute, Aalto Triton, and CSC) and access unique research datasets
- Engage in a vibrant AI + X interdisciplinary research environment at the ELLIS Institute Finland and Aalto University
- Prepare research results for submission to a top-tier AI or interdisciplinary venue

Requirements

- Strong interest in AI for healthcare and/or AI for science
- Strong programming skills in Python and PyTorch
- Solid background in machine learning and deep learning
- Prior experience with generative models (especially diffusion models) is highly preferred
- Prior research or publication experience is a plus

References

[1] About the PI: https://jiancheng-yang.com/
[2] LeFusion: Controllable Pathology Synthesis via Lesion-Focused Diffusion Models. ICLR 2025 (Spotlight)
[3] DiffAtlas: GenAI-fying Atlas Segmentation via Image-Mask Diffusion. MICCAI 2025 (Spotlight)

2103 - Development of Advantage-Based Reinforcement Learning Algorithms

Field of study:	Reinforcement Learning	
For students currently studying:	Master's or PhD	
Number of positions offered:	1-2	
School:	School of Electrical Engineering	
Department:	Department of Electrical Engineering and Automation	
Professor:	Joni Pajarinen	joni.pajarinen@aalto.fi
Academic contact person:	Arsenii Mustafin	arsenii.mustafin@aalto.fi

Q-learning is a well-known reinforcement learning algorithm and was the first one which, combined with neural networks, showed success in complex environments. At the same time, Q-learning tends to be unstable and to overestimate state values. Research has tried multiple ways to address this issue, and while the proposed techniques alleviate the problem somewhat, none of these approaches has fully resolved it.

In this project, we will try to approach the problem using advantage-based algorithms, which are inspired by the geometric interpretation of MDPs and the Safe Reward Balancing [1] algorithm recently developed by Dr. Arsenii Mustafin.

Required background:
The project requires a solid understanding of Q-learning–based methods (Double Q-learning, Dueling Q-learning, IQL, etc.), from both theoretical and practical perspectives, including deep learning engineering.

Reference:
[1] Mustafin, A., Pakharev, A., Olshevsky, A., & Paschalidis, I. (2025). MDP Geometry, Normalization and Reward Balancing Solvers. In International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 2476-2484.

2201 - Applications of quantum computing in optimization

Field of study:	Quantum algorithms	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1	
School:	School of Electrical Engineering	
Department:	Department of Electronics and Nanoengineering	
Professor:	Ilkka Tittonen	ilkka.tittonen@aalto.fi
Academic contact person:	Matti Raasakka	matti.raasakka@aalto.fi

One of the most important near-term applications of quantum computers is optimization. Several different methods and algorithms for optimization exist from quantum annealing to hybrid quantum-classical algorithms. In our research group we study and develop further these methods, as well as test their performance in various small-scale examples of real-world optimization problems. Suitable problems arise in a wide variety of different fields such as telecommunications, bioinformatics, cryptography etc.

Your task during this project is to provide implementations of quantum optimization algorithms and benchmark their performance. The applicant is required (at minimum) to have basic background in quantum computing at the level of an introductory course. Background in machine learning is a definite advantage. The tasks can be adjusted according to the background and the interests of the applicant.

For more information on our research group see https://www.aalto.fi/en/department-of-electronics-and-nanoengineering/micro-and-quantum-systems

2301 - Design and Implementation of an Administrative Web Front-End for AI Summarization

Field of study:	AI, Interface design, Web application	
For students currently studying:	Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering	
Department:	Department of Information and Communications Engineering	
Professor:	Tom Bäckström	tom.backstrom@aalto.fi
Academic contact person:	Lucy Truong	lucy.truong@aalto.fi

In the Trust-M project (https://trustmproject.aalto.fi/), we are studying the trust that migrants place in digital services provided by the City of Espoo. In particular, we have developed a tool for summarizing interactions between migrants and city service personnel. It is a web-front end that records the discussion, passes it to an automatic speech recognizer, and uses AI to summarize it. The focus is on building trustworthy systems and thus all components are self-hosted.

An added functionality that we need is an administrative view, where city personnel can do statistical analysis of the types of services requested, types of migrants, etc. The task includes designing the front-end (e.g. Figma prototyping), implementing the web-based front-end and interfacing with the existing database.

Applicants can be either

1. more design-based such that focus is on designing a good interface,
2. more implementation based, focusing on implementing the web-application and interfacing the database, or
3. any combination of these two.

2302 - Simulating how people read with reinforcement learning

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Bai Yunpeng	byp19971001@gmail.com

Reading is a fundamental cognitive activity that we perform constantly in daily life. Yet, despite its ubiquity, the mechanisms by which humans move their eyes across text, allocate attention, and form comprehension remain only partially understood. Recent advances in computationally rational modeling—and especially reinforcement-learning–based generative reading agents—have shown promising ability to simulate human reading behavior under various conditions, such as natural reading or time pressure.

However, **one of the largest open challenges** is capturing **individual differences**. Native and non-native speakers exhibit markedly different gaze patterns; readers with cognitive or linguistic impairments show characteristic deviations from typical reading; and even within the same population, individuals vary substantially in reading speed, fixation strategies, and comprehension behavior. Emerging large-scale eye-tracking datasets (e.g., OneStop, ZuCo, CELER) provide new opportunities to model these nuanced human differences rather than only normative, averaged behavior.

In this project, we extend a state-of-the-art generative reading model into a **dynamic, individualized simulator** that can reproduce and predict diverse human reading behaviors. This work opens a new direction for modeling not only *how people should read* under optimality assumptions, but also *how different people actually read*—a key capability for education, assistive technology, and human-AI interaction.

The summer intern will contribute to:

- building and refining the reading simulator,
- training reinforcement-learning models,
- evaluating model-human alignment using large eye-tracking datasets,
- running computational experiments, and
- conducting literature reviews in cognitive modeling and machine learning.

Required skills:

- Strong programming ability (Python)
- Background in machine learning, with preference for reinforcement learning
- Passion for understanding human behavior and patience for iterative model development and discussion

References:

[1] Oulasvirta, A., Jokinen, J. P., & Howes, A. (2022, April). Computational rationality as a theory of interaction. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (pp. 1-14).
[2] Bai, Y., Jin, X., Zhao, S., Oulasvirta, A. (2025, December). A resource-rational mechanism for reading.

2303 - AI alignment by combining different forms of Human Feedback

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Jan Kompatscher	jan.kompatscher@aalto.fi

AI alignment is an essential research field that ensures that AI models act in concordance with the needs and wants of human beings.
To train an AI agent to perform a task with reinforcement learning, we need to give it the right numerical reward that accurately attributes a reward/penalty to what the agent should and should not do. We have different approaches for computing such a reward:

1. A mechanistic function defined over the observations and actions of the agent.
2. An artificial neural network that learns to predict the rewards from human preferences.
3. A distance from an optimal behavior policy shown by an expert.

The goal of this research will be to find algorithmic ways for combining these (and other) approaches to better optimize the agent's policy towards the real human intent.

We are collaborating with experts in machine learning and AI from the NII in Tokyo and with an explainable AI lab at ETH in Zürich. We are aiming for top conferences in the ML field like NeurIPS, AAAI, IJCAI, and others.

The summer interns' tasks will be to implement training paradigms for RL models, conduct and evaluate tests, and conduct literature research.

Required skills:

- Programming skills (especially Python), show us your best projects
- Knowledge in machine learning, with a preference for reinforcement learning
- Research skills

References:

[1] Metz, Y., Lindner, D., Baur, R., Keim, D., & El-Assady, M. (2023). RLHF-Blender: A Configurable Interactive Interface for Learning from Diverse Human Feedback (No. arXiv:2308.04332). arXiv. https://doi.org/10.48550/arXiv.2308.04332
[2] Xue, W., An, B., Yan, S., & Xu, Z. (n.d.). Reinforcement Learning from Diverse Human Preferences. Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24). https://www.ijcai.org/proceedings/2024/0586.pdf
[3] Li, J., Luo, B., Xu, X., & Huang, T. (2025). Offline reward shaping with scaling human preference feedback for deep reinforcement learning. Neural Networks, 181, 106848. https://doi.org/10.1016/j.neunet.2024.106848
[4] Kompatscher, J., Shi, D., Varni, G., Weinkauf, T., & Oulasvirta, A. (2025). Interactive Groupwise Comparison for Reinforcement Learning from Human Feedback. Computer Graphics Forum, e70290. https://doi.org/10.1111/cgf.70290

2304 - Uncertainty-Aware Meta-Learning for User Modelling

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Xinyi Wen	xinyi.wen@aalto.fi

User modelling is fundamental to human-AI interaction, enabling systems to understand, predict, and adapt to individual behaviour. This project leverages amortized inference, an emerging framework that combines deep learning with probabilistic reasoning, to enable effective user adaptation. This yields several benefits: (1) uncertainty quantification that captures confidence in predictions, (2) meta-learned knowledge that transfers to new problem instances, and (3) rapid online execution without expensive recomputation, making it ideal for real-time interactive systems.

This project aims to develop amortized inference frameworks for user modelling in human-AI interaction contexts, with applications to AI-assisted design and decision support. We will explore methods for incorporating domain knowledge (such as design principles, task constraints, cognitive models, etc.) as priors, automating efficient interaction design and data acquisition processes, and inferring user latent states and behaviours.

As a summer intern, you will design and implement neural networks for probabilistic user modelling, and evaluate these methods in real-world applications. You will have the opportunity to collaborate closely with researchers from the Finnish Center for Artificial Intelligence (FCAI) and ELLIS Institute Finland, and contribute to work targeting top-tier publication venues.

Requirements:

1. Solid knowledge of deep learning
2. Proficiency in Python and PyTorch programming
3. Basic knowledge of / interest in probabilistic modelling and Bayesian inference
4. Interest in human-computer interaction

References:

[1] Huang, D., Wen, X., Bharti, A., Kaski, S., & Acerbi, L. (2025). ALINE: Joint amortization for Bayesian inference and active data acquisition. *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

[2] Moon, H.-S., Oulasvirta, A., & Lee, B. (2023). Amortized inference with user simulations. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.

[3] Kobalczyk, K., & van der Schaar, M. (2025). Towards automated knowledge integration from human-interpretable representations. *The Thirteenth International Conference on Learning Representations*.

2305 - Human-in-the-Loop Bayesian Optimization for Design

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Thomas Langerak	thomas.langerak@aalto.fi

Human designers rarely approach optimization as a fixed problem. In creative and interaction design work, the space of possibilities shifts as designers explore new directions, reinterpret constraints, and refine their goals. Effective human-in-the-loop optimization requires methods that adapt to evolving problem definitions rather than assume a static design space. This project examines **how to extend Bayesian Optimization (BO) to support such dynamic workflows**. Standard BO assumes a fixed design space, which conflicts with HCI scenarios where users add or remove parameters, adjust bounds, or reshape constraints. These changes break core assumptions of Gaussian Processes and create dimensional inconsistencies that conventional models cannot handle.

The project's technical focus is to **develop a BO framework that remains coherent under structural non-stationarity**. You will explore approaches such as **transfer learning** techniques that model each space modification as a sequential task with informative priors, and **marginalization or imputation strategies** that integrate out removed dimensions and quantify uncertainty for missing variables. The aim is to preserve information across changes while keeping optimization efficient.

As a stretch goal, you will **prototype a human-in-the-loop interface** that lets designers fluidly redefine parameters and constraints, and you will run user studies to evaluate how this flexibility influences creative outcomes and designers' sense of control.

Requirements, you don't need to fulfill all requirements to be considered:

- Bachelor or Master degree in CS (HCI preferred)
- Python proficiency **Preferred**: with BoTorch
- **Preferred**: Frontend and visualization skills in D3.js, React.js, Vue.js, or similar
- **Preferred** Understanding of Gaussian Processes, including kernels, mean functions, and variance

Proposed Reads:

- Bai, Tianyi, et al. "Transfer learning for Bayesian optimization: A survey." *arXiv preprint arXiv:2302.05927* (2023).
- Koyama, Yuki, and Masataka Goto. "Bo as assistant: Using bayesian optimization for asynchronously generating design suggestions." *Proceedings of the 35th annual ACM symposium on user interface software and technology*. 2022.

2306 - AI of the Beholder: VLM-based Evaluation of Character Movement Quality

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Nam Hee Gordon Kim	namhee.kim@aalto.fi

From humanoid robots to digital characters, making AI-generated movements *natural* is a crucial requirement. Despite the need, there is no clear, widely accepted metric for motion quality, and researchers still disagree on its definition. Currently available metrics are designed to assess whether generated motions resemble real data, but they do not generally capture the principles that make movement look natural to our eyes. Meanwhile, querying foundation models is an emerging alternative for evaluating synthesized images and sounds, with evidence of replicating human preferences well. We therefore ask whether Vision-Language Models (VLMs) can also serve as reliable judges of character movement quality.

Our primary objective is then to **assess how effectively state-of-the-art VLMs can evaluate AI-generated character animation quality, measuring their agreement with human judgment and identifying their strengths and limitations.** Achieving this will clarify the feasibility of a VLM-based evaluation and establish a concrete direction toward a robust, widely accepted metric for motion quality.

The summer intern will lead this research investigation, interrogating VLMs to understand how they perceive motion and how faithfully they might replicate human perception of motion quality. We will first focus on prompting VLMs to reproduce recorded human preference data from a publicly available dataset. Then, we will explore the extent to which a VLM-based quality metric can be applied, e.g., to align a generative model's output or to provide a basis for a reward function in deep reinforcement learning (DRL).

The project will involve designing and managing large-scale VLM evaluation experiments, analyzing VLM architectures in explaining their behavior, and potentially conducting a user study to validate the VLM-based metric.

A successful candidate must clearly demonstrate the following:

- A strong academic background in a quantitative field, with prior training in generative models and DRL
- Experience in Python programming with conventional ML, graphics, and visualization frameworks, such as PyTorch, ImGui, Matplotlib, and more. High-performance computing (HPC) skills are a huge bonus.
- Intellectual autonomy and curiosity; ideally in prior publication projects, or alternatively in extracurricular pursuits

Key references:

- Tevet G, Raab S, Gordon B, Shafir Y, Cohen-Or D, Bermano AH. Human motion diffusion model. arXiv preprint arXiv:2209.14916. 2022 Sep 29.
- Figueiredo F, Martinelli G, Sousa H, Rodrigues P, Pedrosa F, Ferreira LN. Echoes of Humanity: Exploring the Perceived Humanness of AI Music. arXiv preprint arXiv:2509.25601. 2025 Sep 29.
- Grötschla F, Solak A, Lanzendörfer LA, Wattenhofer R. Benchmarking Music Generation Models and Metrics via Human Preference Studies. InICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2025 Apr 6 (pp. 1-5). IEEE.
- Wang H, Zhu W, Miao L, Xu Y, Gao F, Tian Q, Wang Y. Aligning human motion generation with human perceptions. arXiv preprint arXiv:2407.02272. 2024 Jul 2.

2307 - Bayesian Optimization for Eliciting Individuals' Creative Potential

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Joongi Shin	joongi.shin@aalto.fi

Computationally modeling individuals' creativity is an important subject in HCI. A model with an accurate prediction of creative behaviors enable reproducing creative individuals' unique thinking patterns and help design proactive assistants. For example, computational model of idea generation task would take a person's personality or preferences as input and generate a sequence of ideas, predicting when the person might experience difficulties in generating new ideas. The key to accurate prediction is (a) capturing the mechanism of human creativity and (b) input data that is useful for the model in reproducing individual differences.

This project aim to tackle the challenges in the later aspect, **focusing on optimizing the process of collecting rich information about individuals in reproducing their behaviors.** The majority of existing approaches require effortful processes, requiring people to perform hour-long surveys. It is also problematic that the existing instruments is not informative enough for computational models, where personality is often captured as a rating such as 'openness = 5' (data). To solve this, we will explore how Bayesian optimization can be applied in the inquiry process. Potential scenario is a cognitive language agent (e.g., generative agents implemented with LLM) iteratively generating the next question to inquire about users. With new information about the user, it will update its belief about the user's characteristics. The formed beliefs will be tested with computational models in real time to check their usefulness in reproducing target human behaviors.

The summer intern will help design, implement, and evaluate this cognitive language agent, and be part of writing a scientific paper.

To this end, the following skills are needed:

a. Programming skills (Python)
b. Knowledge in Bayesian optimization
c. Experience in prompt engineering

Related work:

- Kolluri, Akaash, et al. "Finetuning LLMs for Human Behavior Prediction in Social Science Experiments." Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing. 2025.
- Rossetti, Giulio, et al. "Y social: an LLM-powered social media digital twin." arXiv preprint arXiv:2408.00818 (2024).

2308 - Simulating Health Behaviors to Design Interventions Through Reinforcement Learning

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Pragya Singh	pragyas@iiitd.ac.in

Maintaining a healthy lifestyle is important for both physical and mental well-being, yet it has become increasingly difficult in today's fast-paced world. To address this challenge, researchers are exploring mobile-based health interventions designed to promote adherence to positive habits. However, developing such interventions is challenging, as individuals respond differently to the same prompts depending on their personal characteristics and situational context. Existing approaches often fail to account for these individual differences and contextual factors that shape daily health decisions. Machine learning, behavioral theory, and cognitive modeling offer opportunities to design adaptive and personalized systems that can better support healthy behavior across diverse populations.

This project aims to develop computational models of health behavior that integrate behavioral theories [3] with data-driven techniques to predict and understand individual decision-making [1, 2, 4]. The goal is to design reinforcement learning-based intervention systems that can learn optimal strategies for supporting sustainable behavior changes over time.

The intern will contribute to building the behavior simulator, training reinforcement learning models, and designing the intervention environment.

Required skills:

- Proficiency in Python programming.
- Knowledge of machine learning with a preference for reinforcement learning methods.
- An interest in behavioral science, cognitive modeling, and digital health.

References:

[1] Peng Liao, Kristjan Greenewald, Predrag Klasnja, and Susan Murphy. 2020. Personalized HeartSteps: A Reinforcement Learning Algorithm for Optimizing Physical Activity. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 4, 1, Article 18 (March 2020), 22 pages. https://doi.org/10.1145/3381007

[2] Jiayi Eurus Zhang, Bernhard Hilpert, Joost Broekens, and Jussi P. P. Jokinen. 2024. Simulating Emotions With an Integrated Computational Model of Appraisal and Reinforcement Learning. In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 703, 1–12. https://doi.org/10.1145/3613904.3641908

[3] BJ Fogg. 2009. A behavior model for persuasive design. In Proceedings of the 4th International Conference on Persuasive Technology (Persuasive '09). Association for Computing Machinery, New York, NY, USA, Article 40, 1–7. https://doi.org/10.1145/1541948.1541999

[4] Abdulhai, M., Cheng, R., Clay, D., Althoff, T., Levine, S., & Jaques, N. (2025). Consistently Simulating Human Personas with Multi-Turn Reinforcement Learning. arXiv preprint arXiv:2511.00222.

2309 - UniHand: Unified Hand Control

Field of study:	Artificial intelligence, human-computer interaction	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Electrical Engineering and ELLIS Institute Finland	
Department:	Department of Information and Communications Engineering	
Professor:	Antti Oulasvirta	antti.oulasvirta@aalto.fi
Academic contact person:	Xuejing Luo	xuejing.luo@aalto.fi

We are launching a research project on **Unified Hand Control**, aiming to solve a fundamental challenge in VR/AR and robotic teleoperation: real-world hand tracking is often incomplete. Fingers disappear due to occlusion, joint estimates fluctuate under noise, and some tracking systems provide only partial hand information. Existing methods cannot generate natural or stable hand motions under such imperfect input, nor can they support multiple interaction types within a single framework.

Our goal is to develop a unified model that can take **incomplete, noisy, or partially missing hand-tracking observations** and still produce **physically consistent, task-aware** hand motions. Rather than attempting to cover every possible hand gesture, we focus on three **functional categories** that represent the core patterns of hand–object interaction:

- **Grasping**: forming stable contact by closing the fingers around an object;
- **In-hand Manipulation**: adjusting an object's pose while maintaining a stable grasp;
- **Non-prehensile Interaction**: performing actions such as pressing, pushing, or tapping without grasping.

The goal is for the model to learn the **functional motion patterns** behind these tasks—enabling generalized, robust single-hand control without designing separate controllers for each task.

We are seeking interns with skills or interest in:

- Deep learning (PyTorch), or 3D hand/human modeling;
- Transformers, generative models, or sequence models;
- Experience with VR/AR hand tracking or robotic manipulation is a plus;
- Strong motivation to build multi-task, robust hand-interaction models.

Join us to advance the foundations of next-generation interaction systems.

School of Engineering

Department of Civil Engineering

3201 - EcoNORDICS: A focus on developing Science-based, Eco-Construction Practices

Field of study:	Civil Engineering/Hygrothermal Building Physics	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	3	
School:	School Of Engineering	
Department:	Department of Civil Engineering	
Professor:	Magda Posani	magda.posani@aalto.fi
Academic contact person:	Magda Posani	magda.posani@aalto.fi

The Eco-NORDICS internship positions aim to develop and promote sustainable and locally sourced bio-insulation solutions for deep renovations of the Nordic building stock, with a focus on the Finnish context. The interns will join the ABP (Aalto Building Physics Lab) to learn about locally available, bio-based insulation materials, as well as how to produce and develop them. The three invited students will study the hygrothermal performance and moisture safety of proposed bio-based reno-vation solutions, under present and future climate conditions. The expected findings of this research will include best practice recommendations to make renovations more climate-resilient and climate-neutral.

Position 1) Knowledge development on bio-insulation: Conventional insulation materials like mineral wool and EPS are well-documented for their hygrothermal properties. In contrast, bio-based materials show greater variability and there is a scarcity of existing data. This lack of information complicates reliable predictions for their thermal performance and moisture safety. Thus, acquiring further detailed knowledge about the hygrothermal behaviour of bio-insulation materials is urgently needed. **Objective of the research:** To establish a comprehensive benchmark of hygrothermal data for bio-insulation materials, thus helping support numerical simulations and enriching the scientific literature with new insights and knowledge. **Methods:** literature review, experimental tests, and applied Multiphysics/AI modelling. **Preferred background: Bachelor student from Engineering studies, ideally with previous involvement in projects related to Multiphysics/AI modelling, andPython coding. Previous experience with experimental setups and lab measurements is appreciated but not necessary.**

Position 2) Resilient bio-renovations: The capillary-active and hygroscopic nature of bio-insulation materials can help reduce moisture-related problems exacerbated by climate change in Finland. More guidance needs to be developed to determine the correct design strategies that will support their implementation for climate-resilient renovation solutions. **Objective of the research:** To establish the risks of moisture-related degradation (including mould) and loss of thermal performance posed by climate change to the building envelopes in Finland. To define best practice recommendations for bio-renovations in the Finnish context, in both present and future climate scenarios. **Methods:** Multiphysics modelling and inspections/measurements in existing buildings. **Preferred background: Master student from Engineering studies, ideally with previous involvement in projects related to numerical modelling for buildings/HVAC systems/mould growth and degradation. Previous experience with experimental setups and lab measurements, particularly with biobased building materials, and inspections on-site in existing buildings are appreciated.**

Position 3) Healthier renovations: Biobased materials offer us a unique benefit for indoor comfort and wellbeing in buildings spaces. They are typically characterised by high moisture buffering performance, as so they have the ability to passively regulate indoor humidity levels. By providing moderate humidity levels, they reduce the spread of viruses and bacteria, as well as reduce respiratory infections. This benefit is currently overlooked by building designers, with part of the problem being the lack of reliable simulation tools. **Objective of the research:** Develop a Multiphysics simulation tool able to grasp the moisture buffering value of building materials and the effect of air velocity on their hygrothermal behaviour. **Preferred background: PhD student from Engineering with a specialisation in Building Physics. Previous experience with COMSOL Multiphysics, Computational fluid dynamics, lab measurements and their use in validation/calibration of nu-merical simulation models. Previous experience with bio-based materials (e.g., thermal and acoustic performance) and indoor comfort assessment (on-site measurements and numerical analyses) is appreciated.**

3202 - Digital Construction- Computer Vision perspective

Field of study:	Computer Vision/Geoinformatics	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School Of Engineering	
Department:	Department of Civil Engineering	
Professor:	Olli Seppänen	olli.seppanen@aalto.fi
Academic contact person:	Olli Seppänen	olli.seppanen@aalto.fi

Description: Digital twins, an emerging concept in manufacturing and construction, provide three-dimensional replicas of real structures that facilitate planning, design, construction, and facility management. Precise digital twins have the potential to optimize construction processes, thereby minimizing delays and cost overruns. Achieving this level of accuracy involves automating the capture and analysis of video or point cloud data. However, current techniques in construction require additional research before they can fully convert site data into digital formats information. This project seeks to apply these methods to real construction data to generate useful insights for planning and managing construction projects.

The intern will assist in designing and developing the framework for 360-degree video or point cloud processing related to construction data interpretation.

Therefore, the intern is expected to have the following skills:

- Good coding skills. Preferred language: Python
- Experience in using deep learning libraries like tensorflow, keras, pytorch
- Good Computer vision/3D reconstruiction knowledge
- Interest in object detection/semantic segmentation/COLMAP
- Good communication skills
- Willing to work in a team

The intern is expected to work at the Department of Civil Engineering.

3203 - Simulation of Construction Processes – Digital Twin Perspective

Field of study:	Construction Management	
For students currently studying:	Master's or PhD	
Number of positions offered:	1	
School:	School of Engineering	
Department:	Department of Civil Engineering	
Professor:	Olli Seppänen	olli.seppanen@aalto.fi
Academic contact person:	Mawara Khan	mawara.khan@aalto.fi

Description: Digital twins in Construction is an emerging concept that has drawn attention in the manufacturing as well as construction industries. The digital twins of buildings are a 3D replica of a real structure, which helps plan, design, construct, and can go to facility management of a building. Accurate digital twins will be helpful in streamlining the process of construction to reduce delays and cost overruns associated with construction projects. Modeling and simulation of scenarios by leveraging real-time data from construction sites serves as the backbone of Digital twins.

The intern will participate in the work package of the project that is related to modeling and simulation. The project aims to provide real-time implementation of site situations to analyze and predict the behaviors and activities of the construction workers.

We expect the following skills from the intern:

- Preferred Knowledge and experience in Agent-Based modeling and simulation
- Preferred Knowledge of AnyLogic: Simulation Modeling Software
- Experience in Java /Python programming
- Creative and intelligent mind to develop unique algorithms.
- Good writing and verbal communication skills
- Interactive and collaborative problem solver

The intern is expected to work at the Department of Civil Engineering.

School of Science

Department of Applied Physics

4101 - Biofabricated mycelium cellulose cellular solids for programmable insulation and structures

Field of study:	Physics and materials science	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	2	
School:	School of Science	
Department:	Department of Applied Physics	
Professor:	Mikko Alava	mikko.alava@aalto.fi
Academic contact person:	Juha Koivisto	juha.koivisto@aalto.fi

Bio based foams developed in earlier work have shown high strength to weight ratio and scalability. This project extends that work by integrating mycelium growth and biofabrication as active structure forming mechanisms. Mycelium networks are combined with cellulose-based foams to create hybrid cellular solids with tunable mechanical, thermal, and acoustic properties.

Position 1 focuses on experimental biofabrication. The student will work on controlled mycelium cultivation, foam processing, and post treatments. Mechanical or thermal measurements are used to connect growth conditions to macroscopic performance.

Position 2 focuses on growth simulations and modeling. The student will study mycelium network formation using simplified growth models such as directed percolation, branching processes, or lattice-based growth rules. These models are used to understand connectivity, anisotropy, and density gradients, and to guide experimental design.

The two positions are closely linked and aim to demonstrate programmable, low energy materials suitable for insulation, panels, or architectural elements by combining experiments with physics-based simulations.

Learning outcomes:

- Experience with mycelium biofabrication or growth modeling.
- Understanding growth driven structure property relations.
- Insight into sustainable materials physics for construction.

Background required:

- Interest in experimental materials or computational physics.
- Studies in physics, materials science, or engineering.
- Basic programming skills are useful for the simulation position.

4102 - Photon transport in fusion plasma boundaries

Field of study:	Plasma physics, nuclear fusion	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Applied Physics	
Professor:	Mathias Groth	mathias.groth@aalto.fi
Academic contact person:	Ray Chandra	ray.chandra@aalto.f

A key challenge for future fusion reactors is managing the intense energy deposition into wall surfaces and ensuring long component lifetimes without compromising on performance. An important ingredient in tackling this challenge is the understanding of the particle and heat transport at the fusion plasma boundary, where extremely complex interactions occur between plasma and neutral particles. Fusion boundary physics is commonly explored through large-scale numerical simulations on high-performance computing clusters, which reveals how plasma–neutral interactions shape particle and heat fluxes to the reactor walls. More recently, the inclusion of light or photons into the mix is found to be important for fusion reactors envisioned for power production. The inclusion of photon transport in existing numerical models is a relatively new and rapidly developing area, offering improved accuracy in modeling the plasma boundary and radiation dynamics. Further development includes exploration into physics that uniquely occurs in fusion plasmas.

We are looking for a candidate to explore one of the two projects listed below:

(1) Applying advanced photon transport models to present-day high-power fusion devices. The student is expected to predict the impact of photon interactions to the description of plasma boundaries in present-day fusion machines using numerical models developed within our group. Suitable for Master level students.

(2) Exploration into ultraviolet-resonant transition lines of molecular deuterium. A central task will be to identify energy level transitions of the deuterium molecule within its rovibrational system using existing energy level data, that closely resembles the main ultraviolet wavelength of fusion plasmas. Suitable for Bachelor and Master level students.

Required background:
The candidate should have basic knowledge of atomic and molecular physics through coursework or self-study and basic experience in programming with any language (e.g. Python). Experience with spectroscopy and/or the Fortran programming language is considered a plus.

4103 - Tensor-Network Tight-binding Hamiltonian for fractals

Field of study:	Theory of quantum materials, quantum many-body physics	
For students currently studying:	Bachelor's or Master's	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Applied Physics	
Professor:	Jose Lado	jose.lado@aalto.fi
Academic contact person:	Anouar Moustaj Tiago Antão	anouar.moustaj@aalto.fi tiago.antao@aalto.fi

Background: Representing large tight-binding Hamiltonians directly in real space becomes infeasible once system sizes reach millions or billions of sites, as encountered in moiré and super-moiré structures. Tensor-network techniques provide a way around this bottleneck by mapping the single-particle Hamiltonian onto an auxiliary many-body pseudospin chain, allowing the Hamiltonian to be encoded as a compact matrix-product operator instead of an enormous sparse matrix. This compression leverages the structured spatial modulations typical of moiré systems, keeping bond dimensions manageable even at extreme sizes. Within this framework, spectral quantities can be computed efficiently through a tensor-network kernel polynomial method. This strategy makes it possible to compute spectral functions, local observables, and even momentum-resolved quantities using a quantum Fourier transform for systems well beyond 10^9 sites, far exceeding the scalability of conventional tight-binding approaches.

Project Expectation: The project will begin with reproducing selected results from arXiv:2503.04373 to gain familiarity with the tensor-network framework. The next step is to construct a tight-binding Hamiltonian on a one-dimensional chain whose onsite or hopping terms are modulated by a function defined on the fractal Cantor discontinuum. Exploiting the hierarchical structure of the Cantor set, a pseudo-spin-1 encoding can be employed to formulate the corresponding fractal Hamiltonian in tensor-network form. Once the model is built, the goal is to compute its spectral properties—such as density of states and local spectral features—and to explore possible extensions to alternative fractal geometries.

References:

- Phys. Rev. Research 7, 043288 (2025), https://journals.aps.org/prresearch/pdf/10.1103/krjp-mn4v
- arXiv:2506.05230 (2025), https://arxiv.org/pdf/2506.05230

Prerequisites:

- Quantum Mechanics
- Tight-Binding models
- Many-body Physics
- Tensor Networks

4104 - Radiation damage modelled with two-temperature molecular dynamics

Field of study:	Physics	
For students currently studying:	Master's	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Applied Physics	
Professor:	Andrea Sand	andrea.sand@aalto.fi
Academic contact person:	Andrea Sand	andrea.sand@aalto.fi

Description: Particle irradiation modifies the physical and mechanical properties of materials, and plays an increasing role in modern technological developments. For example, climate change is driving the need for green energy, with nuclear fusion and next generation fission standing as two of the strongest candidates for efficient and reliable energy production of the future, yet the challenges posed to reactor materials in the high radiation environments are significant. Modelling provides an essential tool for predicting the response of reactor components in future nuclear devices. The damage in materials created by energetic impacting particles is highly sensitive to the mechanisms of dissipation of the impinging particle's kinetic energy.

This summer project involves performing simulations employing a recently developed atomistic model, which accounts for energy dissipation in unprecedented detail, to predict the primary radiation damage in fusion-relevant structural materials under different incident neutron and ion energies. Focus will be on analysis of the surviving damage, including defect numbers and morphology, and comparison to experiments. The student will gain knowledge of the processes of radiation damage formation in metals, learn the basics of performing molecular dynamics simulations of highly non-equilibrium events, and develop a familiarity with high performance computing environments.

Necessary skills:

- Experience in programming, e.g. with Python, is highly desirable.
- The candidate should also have basic knowledge of solid-state physics and computational physics.
- Previous experience of molecular dynamics or high-performance computing is considered a plus.

4105 - Dark optical-Paul trapping of large mass particles in ultra-high vacuum / Entangled light-matter states in polariton optomechanical systems

Field of study:	Quantum physics	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Applied Physics	
Professor:	Anton Zasedatelev	anton.zasedatelev@aalto.fi
Academic contact person:	Anton Zasedatelev	anton.zasedatelev@aalto.fi

The Macroscopic Quantum Optics group offers 1 internship position (experimental) to join one of the two ongoing projects:

(1) Dark optical-Paul trapping of large mass particles in ultra-high vacuum.

(2) Entangled light-matter states in polariton optomechanical systems.

We are seeking a passionate and driven student/researcher with a background in at least one of the following areas:

- atomic, molecular, and optical (AMO) physics
- quantum optics
- optomechanics
- nonlinear or ultrafast optics

MQO is uniquely positioned to combine state-of-the-art experimental facilities with top theoretical expertise under one roof.

For more information, please visit our webpage:
https://www.aalto.fi/en/department-of-applied-physics/macroscopic-quantum-optics-mqo

4201 - A Principled Approach to Watermarking Diffusion Language Models

Field of study:	LLM, Diffusion Models, Algorithm, Theory	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Qi Chen	qi.chen.inf@aalto.fi
Academic contact person:	Qi Chen	qi.chen.inf@aalto.fi

This topic investigates how to design reliable, efficient, and theoretically grounded watermarking methods for diffusion-based language models. The project explores how to embed, detect, and evaluate watermarks in generated text while ensuring robustness, minimal quality degradation, and strong guarantees against removal or evasion.

Background:

- Applicants should be comfortable with machine learning fundamentals, probabilistic modeling, and Python-based experimentation.
- Prior experience with generative models, diffusion processes, or trustworthiness topics is beneficial but not strictly required.

Some related works:

- Gloaguen, Thibaud, Robin Staab, Nikola Jovanović, and Martin Vechev. "Watermarking diffusion language models." *arXiv preprint arXiv:2509.24368* (2025).
- Wu, Linyu, Linhao Zhong, Wenjie Qu, Yuexin Li, Yue Liu, Shengfang Zhai, Chunhua Shen, and Jiaheng Zhang. "Dmark: Order-agnostic watermarking for diffusion large language models." *arXiv preprint arXiv:2510.02902* (2025).
- Bagchi, A., Bhimaraju, A., Choraria, M., Alabi, D., & Varshney, L. R. (2025). Watermarking Discrete Diffusion Language Models. *arXiv preprint arXiv:2511.02083*.
- Liu, Y., Li, Z., Backes, M., Shen, Y. and Zhang, Y., 2023. Watermarking diffusion model. *arXiv preprint arXiv:2305.12502*.
- Zhang, H., Edelman, B. L., Francati, D., Venturi, D., Ateniese, G., & Barak, B. (2024, July). Watermarks in the sand: impossibility of strong watermarking for language models. In *Forty-first International Conference on Machine Learning*.
- He, Haiyun, Yepeng Liu, Ziqiao Wang, Yongyi Mao, and Yuheng Bu. "Theoretically Grounded Framework for LLM Watermarking: A Distribution-Adaptive Approach." *arXiv preprint arXiv:2410.02890* (2024).

4202 - Computational Perspectives on Democracy

Field of study:	Computational Social Science, Computational Politics, Computational Legal Studies, Legal Data Science, Computational Social Choice	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1-3	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Corinna Coupette	corinna.coupette@aalto.fi
Academic contact person:	Corinna Coupette	corinna.coupette@aalto.fi

Supported by the ERC Starting Grant *CompLex – Toward a Computational Theory of Legal Complexity*, the Telos Lab at Aalto University invites applications for several projects developing computational perspectives on democracy. While modern societies and technologies have changed dramatically over the past decades, the structure and procedures of our legal systems have remained largely unaltered. In this project, we analyze the institutions and procedures existing legal systems, investigate their shortcomings, and explore potential strategies to improve democratic processes.

Depending on interns' preferences and skills, topics for summer projects include:

1. **Data science for democracy and legal systems.** Collect, preprocess, model, and analyze legal or political data, develop methods to study domain-specific data across multiple modalities, including networks and texts.
   Desirable background: Experience handling complex (network and text) data, especially from law or political science.

2. **Agent-based modeling of policymaking systems.** Build simulations of legal systems or policymaking processes to evaluate their performance and compare them with state-of-the-art alternatives.
   Desirable background: Experience with computational modeling and simulation design.

3. **Theoretical approaches to democratic processes.** Analyze the mathematical and computational properties of existing procedures for collective decision-making, e.g., regarding their power dynamics, fairness, and transparency, and investigate the properties of promising alternatives.
   Desirable background: Experience with mathematical modeling and interest in societal applications of theoretical computer science.

For an overview of the type of work conducted in the lab, see https://www.coupette.io/publications/.

Programming languages: Python

Libraries/frameworks for handling: dataframes, databases, network modeling and analysis, computational modeling and simulation, data analysis, data visualization, data mining, machine learning, natural language processing, web scraping.

Potential tasks (with different emphases in different subprojects): literature review; data collection, data preprocessing, data modelling, and data engineering; computational modeling and simulation; problem formalization and theoretical analysis; method development and (experimental) validation; data analysis and visualization; software development, validation, and testing; collaboration and community engagement.

4203 - Machine Learning: Foundations and New Frontiers

Field of study:	All areas of Machine Learning, including, but not restricted to Generative Models, Geometric Deep learning, LLMs, Neural ODEs, Physics-informed Learning, Computer Vision, Computational Biology, and (Applied) Math.	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	5-10	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Vikas Garg	vgarg@csail.mit.edu
Academic contact person:	Vikas Garg	vikas.garg@aalto.fi

Applications are invited for various internship positions in our group, see e.g., [6]-[16] below about our contributions to deep learning, drug discovery, graph neural networks, generative models, climate modelling etc. An ideal student would be eager to push the frontiers of science; have strong mathematical, theoretical, statistical, or algorithmic background; and be comfortable with programming in Deep Learning. **We particularly invite students with strong Math, LLM, Bioinformatics, and Physics backgrounds to apply.** We value diversity and encourage candidates from underrepresented backgrounds to apply. The positions in our group are extremely competitive (~1700 intern applications last year) but very rewarding (see, e.g., https://shorturl.at/SGagV and https://shorturl.at/4DZ4Y for experiences shared by former members).

Interns in our group from previous years have produced stellar research [1-5], including, a publication accepted with Oral presentation at NeurIPS 2023 [1].

Topics of particular interest include but are not limited to:
(1) Generative Models, Neural ODEs/PDEs/SDEs, Deep Equilibrium Models, Implicit Models
(2) Large Language Models, Reinforcement Learning, Agentic AI, and AI-assisted human-guided models
(3) Automated theorem proving and neurosymbolic reasoning
(4) Protein Design, Material Design, Molecular Simulations, and Drug Discovery
(5) (Temporal) Graph Neural Networks, Topological Deep Learning, Topological Data Analysis (e.g., Persistent Homology)
(6) Differential Geometry/Information Geometry/Algebraic/Spectral Methods for Deep Learning
(7) (Approximate) Equivariant and Invariant models
(8) Fair, interpretable, and explainable methods

Publications by our former interns:
[1] J. Immonen(*), A. Souza (*), and V. Garg. Going beyond persistent homology using persistent homology. NeurIPS (2023).
[2] K. Brilliantov, A. Souza, and V. Garg. Compositional PAC-Bayes. NeurIPS (2024).
[3] T. Pham and V. Garg. What do Graph Neural Networks learn? Insights from Tropical Geometry. NeurIPS (2024).
[4] M. Ji, A. Souza, and V. Garg. Graph Persistence goes Spectral. NeurIPS (2025).
[5] M. Ji, A. Souza, and V. Garg. On topological descriptors for graph products. NeurIPS (2025).

Representative publications by our group (https://people.csail.mit.edu/vgarg/):
[6] R. Karczewski et al. Devil is in the Details: Density Guidance for Detail-Aware Generation with Flow Models. ICML (2025).
[7] Laabid(+), Rissanen(+) et al. Equivariant Denoisers Cannot Copy Graphs: Align your Graph Diffusion Models. ICLR (2025).
[8] J. Ingraham, V. Garg, R. Barzilay, and T. Jaakkola. Generative Models for Protein Design. NeurIPS (2019).
[9] V. Garg, S. Jegelka, and T. Jaakkola. Generalization and Representational Limits of Graph Neural Networks. ICML (2020).
[10] V. Garg and T. Jaakkola. Solving graph compression via Optimal Transport. NeurIPS (2019).
[11] T. Garipov et al. Compositional Sculpting of Iterative Generative Processes. NeurIPS (2023).
[12] Y. Verma, M. Heinonen, and V. Garg. AbODE: Ab initio antibody design using conjoined ODEs, ICML (2023).
[13] A. Souza, D. Mesquita, S. Kaski, and V. Garg. Provably expressive temporal graph networks. NeurIPS (2022).
[14] G. Mercatali, A. Freitas, and V. Garg. Symmetry induced disentanglement on graphs. NeurIPS (2022).
[15] Karczewski et al. What Ails Generative Structure-based Drug Design. AISTATS (2025)
[16] Y. Verma et al. Climate and Weather Forecasting with Physics-informed Neural ODEs. ICLR (2024)

4204 - Social Intrinsic Motivation

Field of study:	Computer Science / Artificial Intelligence	
For students currently studying:	Master's or PhD	
Number of positions offered:	1-2	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Christian Guckelsberger	christian.guckelsberger@aalto.fi
Academic contact person:	Christian Guckelsberger	christian.guckelsberger@aalto.fi

Our goal is to **advance social AI**, i.e. artificial agents capable of supporting us proactively or reactively in a range of virtual or embodied tasks. However, existing social AI and artificial agents more generally still struggle in delivering reliable support in situations that were not anticipated at design-time, and their decision-making processes may be hard to understand by interaction partners or operators. This project seeks to tackle these challenges by advancing **Social Intrinsic Motivation**, i.e. computational models of intrinsic motivation applied to the interaction of multiple agents. **Intrinsic motivation (IM)**, with curiosity being a well-known example, is crucial to human development, allowing us to acquire new knowledge and skills in the absence of specific tasks, and to adapt to new situations. Computational models of IM inherit these properties, making artificial agents more independent of external instruction and allowing them to perform better in previously unknown tasks. Typically formalised in the framework of **reinforcement learning**, computational models of IM are not black boxes, which enables explanation of the emerging behaviour and supports trust by the AI's designers and interaction partners.

Social Intrinsic Motivation seeks to leverage the power of IM in the interaction of multiple agents. One example could be motivating an AI to optimise their partner's intrinsic motivation, thus empowering them through its own behaviour. Pioneering work led to **exciting findings** such as emergent support of human interaction partners across tasks and without external instruction [1-4] or considerably improved cooperation within societies of artificial agents [5,6]. **Projects under this heading can take various forms**. Interns may support research on developing a better understanding of the space of possible SIM, the systematic development of new models informed by psychology, or evaluation in large-scale simulation or smaller-scale user studies with people.

The Autotelic Interaction Research (AIR, www.autotelic.science) group at Aalto supports this work through a **track record of research** on computational IM, Social Intrinsic Motivation, and translational research with psychology.

Successful applicants must demonstrate a background in computer/data science (at least far into their Master's), and good knowledge (courses, project work) of reinforcement learning. Knowledge of computational IM is a merit.

References:

[1] Guckelsberger, C., Salge, C. and Colton, S., 2016. Intrinsically motivated general companion npcs via coupled empowerment maximisation. In Proc. IEEE Conference on Computational Intelligence and Games (CIG).

[2] Guckelsberger, C., Salge, C. and Togelius, J., 2018. New and surprising ways to be mean. In Proc. IEEE Conference on Computational Intelligence and Games (CIG).

[3] Guckelsberger, C., "Coupled Empowerment Maximisation (CEM)". In: Guckelsberger, C., 2020. Intrinsic Motivation in Computational Creativity Applied to Videogames. PhD Thesis. Queen Mary University of London.

[4] Salge, C. and Polani, D., 2017. Empowerment as replacement for the three laws of robotics. Frontiers in Robotics and AI, 4.

[5] van der Heiden, T., van Hoof, H., Gavves, E. and Salge, C., 2022. Reliably Re-Acting to Partner's Actions with the Social Intrinsic Motivation of Transfer Empowerment. In Proc. ALife, 2022.

[6] Jaques, N., Lazaridou, A., Hughes, E., Gulcehre, C., Ortega, P., Strouse, D.J., Leibo, J.Z. and De Freitas, N., 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In Proc. ICML.

4205 - Bayesian physician — training a Bayesian neural network to emulate physician's decisions

Field of study:	Bayesian neural networks	
For students currently studying:	Master's	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Samuel Kaski	
Academic contact person:	Jorge Loría	jorge.loria@aalto.fi

Doctors decide treatments and diagnoses in varied manners for their patients. This project aims to fit a Bayesian neural network (BNN) that can predict the treatments and diagnoses each doctor would assign to different patients. Understanding the variability of decisions between different doctors with similar patients is a task that will be studied within this project. Additionally, the obtained model will be used for a downstream task. The importance of using a Bayesian approach is that it allows to simulate from the posterior of models conditional on the observations, which is key for other downstream tasks such as fitting foundational models.

The ability to handle large volumes of data in a limited computational environment is expected.

Prerequisites:

- torch
- variational inference
- Bayesian methods
- knowledge of causal inference is desirable but not expected

References:

- Zhang, C., Bütepage, J., Kjellström, H., and Mandt, S. (2018). Advances in Variational Inference. IEEE Transactions on Pattern Analysis and Machine Intelligence, 41(8):2008–2026.
- Liu, Q., & Wang, D. (2016). Stein variational gradient descent: A general purpose Bayesian inference algorithm. *Advances in neural information processing systems*, *29*.
- Papamarkou, T., Skoularidou, M., Palla, K., Aitchison, L., Arbel, J., Dunson, D., ... & Zhang, R. (2024). Position: Bayesian deep learning is needed in the age of large-scale AI. *arXiv preprint arXiv:2402.00809*.
- Hollmann, N., Müller, S., Purucker, L., Krishnakumar, A., Körfer, M., Hoo, S. B., ... & Hutter, F. (2025). Accurate predictions on small data with a tabular foundation model. *Nature*, *637*(8045), 319-326.

4206 - Geometric algorithms in Euclidean and hyperbolic spaces

Field of study:	Theoretical Computer Science	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1-2	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Sándor Kisfaludi-Bak	sandor.kisfaludi-bak@aalto.fi
Academic contact person:	Saeed Odak	saeed.odak@aalto.fi

Our computational geometry research group is looking for a summer intern to help us with our research related to geometric algorithms in Euclidean and hyperbolic spaces. We expect a good understanding of mathematics (e.g. in discrete math, graph theory, geometry), algorithms, and theoretical computer science.

Potential topics include geom. optimization, parameterized geom. algorithms, fine-grained complexity, geom. approximation, spanners, geom. intersection graphs. Each of these is studied in both Euclidean and hyperbolic settings. The position is focused on fundamental theory, we prove theorems about the algorithmic properties of geometric problems. (There is no software engineering involved.) In the past two years, these summer projects have culminated in research articles co-authored with the interns.

4207 - Probabilistic machine learning and generative AI

Field of study:	Machine learning	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	2	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Arno Solin	arno.solin@aalto.fi
Academic contact person:	Mohammad Vali	mohammad.vali@aalto.fi

We are looking for motivated interns to join ongoing research projects in probabilistic machine learning, with opportunities across:

(1) structured and efficient inference and generative modelling (mostly diffusion models),
(2) uncertainty quantification (UQ) in deep learning (including large-scale models), and
(3) multi-modal modelling (mostly VLMs), including semantic understanding + 3D scene reconstruction (Gaussian splatting).

You will work on cutting-edge research problems such as uncertainty-aware neural networks, approximate inference methods, and probabilistic approaches to vision–language and 3D reconstruction. The projects are part of a broader initiative to advance the foundations and practice of modern machine learning.

A strong candidate typically has:

(1) Background in probabilistic modelling and approximate inference;
(2) solid understanding of machine learning and deep learning;
(3) strong programming skills in Python and experience with at least one of PyTorch/JAX; and most importantly (4) curiosity, creativity, and a collaborative mindset.

In your application, briefly highlight:

(1) Your relevant coursework and/or research experience;
(2) the topics you are most excited about (UQ, vision–language, 3D, etc.); and
(3) links to code, papers, or projects (if available)

For representative publications and background, see the supervisor's homepage: https://arno.solin.fi

4208 - Distributed quantum computing

Field of study:	Theoretical computer science	
For students currently studying:	Bachelor’s, Master’s, or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Jukka Suomela	jukka.suomela@aalto.fi
Academic contact person:	Jukka Suomela	jukka.suomela@aalto.fi

Our research group "Distributed Algorithms" is looking for a summer intern to help us with our research related to distributed quantum algorithms. Our work is at the intersection of mathematics (especially discrete mathematics and graph theory), theoretical computer science (especially theory of distributed computing and distributed graph algorithms), and quantum physics (especially quantum information theory).

For more information on our research and our research group, see https://research.cs.aalto.fi/da/

4209 - Massively Parallel Algorithms for Graph Problems

Field of study:	Theoretical Computer Science	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Jara Uitto	jara.uitto@aalto.fi
Academic contact person:	Jara Uitto	jara.uitto@aalto.fi

Parallel processing of data and distributed computing are gaining attention and becoming increasingly vital as the data sets and networks we want to process are overgrowing the capacity of single processors. To understand the potential of modern parallel computing platforms, many mathematical models have emerged to study the theoretical foundations of parallel and distributed computing. In this project, we study algorithm design in these models with a particular focus on the Massively Parallel Computing (MPC) and Local Computation Algorithms (LCA) models.

The problems we study are often in (but not limited to) the domain of graphs, that serve as a very flexible representation of data. We are interested in, for example, the computational complexities of classic problems such as finding large independent sets, matchings, flows, clustering problems, etc.

Required background:
The applicant is assumed to have a solid knowledge of mathematics, knowledge on the basics of graph theory, and a good command of English. No prior knowledge in distributed computing is required, although it might be helpful.

4210 - Bayesian workflow

Field of study:	Bayesian computational modeling	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Aki Vehtari	aki.vehtari@aalto.fi
Academic contact person:	Osvaldo Martin	osvaldo.martin@aalto.fi

You will take part in developing computational diagnostic tools for different parts of Bayesian workflow (see, e.g. https://arxiv.org/abs/2011.01808). Possible more specific topics include model checking diagnostics, cross-validation, better priors, inference diagnostics.

Prerequisites: Bayesian inference and MCMC

4211 - The psychology of Human-AI interaction

Field of study:	Human-Computer Interaction, AI, Psychology	
For students currently studying:	Bachelor's, Master's, or PhD	
Number of positions offered:	1-2	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Robin Welsch	robin.welsch@aalto.fi
Academic contact person:	Robin Welsch	robin.welsch@aalto.fi

Artificial intelligence (AI) technologies fundamentally affect our every-day lives. Human-centred AI technologies are designed to support humans and extend human capabilities, raising high expectations concerning our AI-augmented abilities. This research internship will run empirical studies to investigate how AI systems increase cognitive performance.

The main task in this internship are the design, planning, and execution of an empirical user study. This includes the creation of prototypes, the preparation of study materials but also doing research with users in the lab. For a testimonial of last year's interns see: http://human-ai-interaction.com/allgemein/summer-is-over-and-our-foreign-students-are-leaving/

Candidates for the intern position should have the following skills:

- knowledge in quantitative user studies and/or AI
- good knowledge of quantitative data analysis
- strong interest in experimental psychology research

You will gain experience in the following areas:

- Artificial intelligence and mental models
- Usability & User Experience
- Conducting User studies
- Application of psychology to human-computer interaction
- Artificial Intelligence

4212 - Designing New Intimate Wearables for Reproductive Health

Field of study:	Human Computer Interaction, Interaction Design	
For students currently studying:	Master's or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Deepika Yadav	deepika.yadav@aalto.fi
Academic contact person:	Deepika Yadav	deepika.yadav@aalto.fi

As personal health technologies are becoming increasingly advanced and intimate, they open pathways for more expressive and body-aligned forms of care that can be worn, felt, and embodied in daily routines. Yet, current approaches often overlook the significance of integrating the tactile, sensory, and aesthetic dimensions of how we experience our body when interacting with or wearing these technologies. The project aims at designing new intimate wearables that use postures, movements, and bodily rhythms to offer seamless interaction possibilities for improving bodily expression, awareness and deeper connection with the body.

In particular, the focus will on bodily changes that occur across reproductive life stages —including menstruation, pregnancy, postpartum, breastfeeding, and menopause—which bring complex transformations in the health and wellbeing of an individual. These changes are under-supported and under designed for, which limits how we understand and connect with our bodies. Within HCI and Interaction Design, emerging works in intimate care and reproductive health has begun to push beyond quantified and medicalized approaches, towards integrating design that are more attuned to our relationships with our bodies such as by embracing touch of the cervical mucous through a finger-worn sensor [1], improving awareness of the pelvic floor anatomy and its functions through a soft-robotics chair that touches the pelvic floor and areas surrounding the intimate [3], and generatively designing for menstrual pain [2],. These have contributed compelling design exemplars, case studies, analysis and conceptual framings that reveal the generativity of designing for—and with—the body.

In this project, the student will design and develop prototypes of **soft wearable technologies suitable for use as intimate wear or garments—**creating seamless, body-aligned experiences that support new forms of bodily interaction. The focus will be on enabling extended wear by harvesting bodily movements, postures, rhythms, or other physical changes so that users can incorporate the wearable into daily life. The work will involve exploring materials, textiles, and methods for embedding electronics for sensing and feedback that are lightweight and noninvasive, supporting intimate, slow, and subtle forms of interaction.

The thesis will contribute to emerging research in soma design [4] and interaction design, in application areas concerned with design, health, body, e-textiles, and wearables.

The project is ideal for: candidates with strong prototyping skills in wearable, e-textiles or soft-robotics, and an interest in Interaction Design research for health.

References:

1. Campo Woytuk, Nadia, et al. "Toward Feminist Ways of Sensing the Menstruating Body." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 2025.
2. Park, Joo Young, et al. "Designing Touch Technologies for and with Bodies in Menstrual Discomfort." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 2025.
3. Yadav, Deepika, et al. "A Route to Somatic Literacy of the Pelvic Floor through Technology-Initiated Touch." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 2025.
4. Hook, Kristina. *Designing with the body: Somaesthetic interaction design*. MIt Press, 2018.

4213 - Analysing Network Support for Women in Scientific Careers

Field of study:	Computational Social Science, Network Science, Human-Computer Interaction	
For students currently studying:	Master's	
Number of positions offered:	2	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Barbara Keller Deepika Yadav	barbara.keller@aalto.fi deepika.yadav@aalto.fi
Academic contact person:	Anh-Duong Nguyen Deepika Yadav	anh-duong.nguyen@aalto.fi deepika.yadav@aalto.fi

The under-representation of Women in academia is often explained among other factors, via the concept of a "leaky pipepline", where a higher share of women workforce opt-out along the different stages from graduation over PhD programmes to tenure-track positions. The lack of social capital, role models, peer support and mentorship is frequently cited as a major driver of these phenomena. From these phenomena, mothers are even more severely affected [3]. We are interested in approaching the leaky pipeline problem from a network science perspective, investigating tie formation mechanisms such as homophily (connecting with similar others), triadic closure (connecting with mutual contacts), and, preferential attachment (connecting with popular individuals), etc. [2, 1]. Depending on our interns' preferences and skills, topics for the summer project can include:

1. Network Models for Motherhood Penalty Develop and/or extend random graph models (Barab´asi–Albert model, configuration model, dynamic model) to analyze the inequality phenomena with mathematical reasoning and numerical simulation. This includes proposing important parameters and identifying network mechanisms that accelerate or modulate inequality based on the experience of women in scientific careers.

2. Gender difference in tie formation mechanisms Investigate real- world social ties and the underlying factors influencing the social ties and examine how they impact women's growth in scientific careers as they navigate different stages in their professional journeys. This will include collecting data from existing records as well as from participants through interviews, surveys, and workshops. Through a mixed-method analysis approach empirical and qualitative insights will be developed on the women's experience about the network support.

We are planning to hire two students for this project. One with a background in quantitative approaches and one with a background in qualitative approaches.

The following skill sets are beneficial: Mathematical analysis and programming experience (preferably in Python), as well as familiarity with graph theory, and social network analysis methods and/or candidates with experience in designing, conducting and analyzing interviews, workshops, and surveys on this topic.

Potential tasks: Literature Review, Questionnaires Design, Data Mining, Data Processing, Statistical Analysis, Thematic Analysis, and Visualization, Mathematical Modeling and Numerical Simulation.

References:

[1] Aili Asikainen et al. "Cumulative effects of triadic closure and homophily in social networks". In: Science Advances 6.19 (May 2020). issn: 2375-2548. doi: 10.1126/sciadv.aax7310.

[2] Chen Avin et al. "Homophily and the Glass Ceiling Effect in Social Networks". In: Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science. ITCS'15. ACM, Jan. 2015, pp. 41–50. doi: 10.1145/2688073.2688097.

[3] Stephen J. Ceci et al. "Women in Academic Science: A Changing Land- scape". In: Psychological Science in the Public Interest 15.3 (Nov. 2014), pp. 75–141. issn: 1539-6053. doi: 10.1177/1529100614541236

4214 - Change in web browsing behavior across shock events

Field of study:	Computational Social Science, Data Science	
For students currently studying:	Master's or PhD	
Number of positions offered:	1	
School:	School of Science	
Department:	Department of Computer Science	
Professor:	Juhi Kulshrestha	juhi.kulshrestha@aalto.fi
Academic contact person:	Juhi Kulshrestha	juhi.kulshrestha@aalto.fi

By analyzing longitudinal, fine-grained web-browsing traces, in this project we will analyze how people's online browsing behavior is affected by exogeneous events. Further, we will examine whether the impact of exogenous shocks is uniform across the whole population or whether individuals with certain behavioral and socio-demographic characteristics are more or less vulnerable to the shocks.

Required skills: good python programming skills, experience with data science methods and statistical models.