Internship Programme

Project list for the calendar year 2026

Department of Mathematics and Mathematical Statistics
Umeå University, Umeå, Sweden

About the programme

Our internship programme is a one-term opportunity based in Umeå, where students work directly with researchers in their chosen field. The programme provides a unique chance to bridge the gap between academic studies and their application in a research environment.

Eligibility and Applications

The programme is primarily aimed at graduate students from our partner universities, but applications from students at other universities are also welcome. Application rounds are held twice a year.

Autumn round: September 15 – October 15

Students apply for **one** specific project that best matches their academic background and research interests. Applications are submitted via our online form and must include:

- Curriculum Vitae (CV)
- Transcript of records (in English)
- Motivation letter

Support Provided

- Coverage of round-trip travel to Umeå (economy class, funded by the department).
- Assistance with securing student accommodation (costs covered by the intern).
- Support in applying for an Erasmus+ stipend through the home university.
- Access to a supportive work environment at Umeå University, including a workspace, necessary materials, and insurance.

Selection Procedure

Applications are reviewed by a working group in collaboration with the prospective supervisors. Candidates are selected based on eligibility, academic record, and compatibility with the chosen project.

Content of the Document

This document provides information about the projects available in the current application round, along with the corresponding supervisors. Applicants must select one project to apply for. The projects are organised by research area and supervisor.

For more information and access to the application form, please visit:

https://www.umu.se/en/department-of-mathematics-and-mathematical-statistics/education/internship-programme/

Discrete Mathematics

Victor Falgas Ravry

Victor Falgas Ravry is an Associate Professor in Mathematics. His main research interests lie in extremal combinatorics and discrete probability — interests which are shared by several others members of the Discrete Mathematics group at Umeå. Typical topics include extremal graph theory, Ramsey theory, games on graphs and random graph models.

Webpage: https://www.umu.se/en/staff/victor-falgas-ravry/

Project 1: Maker–Breaker percolation games

Maker–Breaker games were introduced in an influential paper of Chvátal and Erdős [2] in the late 1970s and form a rich and widely studied class of positional games.

The board in a Maker–Breaker game consists of a set B (for example, the set of edges of a graph, or the cells of a tic-tac-toe grid). Two players, known as Maker and Breaker, take turns claiming elements of the board B. Maker's goal is to claim all elements of some member W of a pre-specified family $\mathcal{W} \subseteq \mathcal{P}(B)$ of winning sets (for example, a winning set could be all the edges of a triangle, or three squares on the same row, column or diagonal of a tic-tac-toe grid). Breaker's aim is to prevent this from happening by claiming at least one element from each $W \in \mathcal{W}$.

As the vague and general description above suggests, there are many games within this class, and many questions one could ask: who has a winning strategy under optimal play? how long will the game last? what happens if one of the players selects his or her moves at random? what happens if we allow one of the players to claim more elements of the board per turn than the other player?

A rich theory has developed in a bid to answer these questions (see e.g. the book of

Beck [1] and the survey of Stojaković and Szabó [8]), and to elucidate some intriguing connections between these games and deep phenomena in discrete probability and Ramsey theory. As Tibor Szabó has pointed out, these delightfully fun and deceptively light-hearted combinatorial games have contributed to many of the most significant advances in discrete mathematics in the past half century, and are well worth studying for their own sake.

A few years ago, I introduced a new class of Maker–Breaker games with motivation coming from percolation theory. These games are played on a (large) graph G with two specified vertices u and v. Maker and Breaker take turns claiming sets of m and b edges of G respectively, and Maker's aim is to build a path joining u to v.

These Maker–Breaker percolation games may be viewed as generalisations of the classical Shannon switching game (which corresponds to the special case $m = b = 1$, and was fully resolved by Lehman [7] in the 1970s), and have been the subject of two articles by my postdoc Nicholas Day and myself [4, 3], and more recently to two lovely papers of Dvořák, Mond and Souza [5]. As can be seen from a cursory glance at the concluding sections of these four papers, a great many open questions remain — many of which would be highly suitable for an internship, including in particular the as-yet unstudied site percolation version of the game (where Maker and Breaker claim vertices rather than edges of G).

References

[1] József Beck. *Combinatorial games: tic-tac-toe theory*, volume 114. Cambridge University Press Cambridge, 2008.

[2] Vašek Chvátal and Paul Erdős. Biased positional games. Pp. 221–229 in *Annals of Discrete Mathematics*, volume 2. Elsevier, 1978.

[3] A. Nicholas Day and Victor Falgas-Ravry. Maker–breaker percolation games I: crossing grids. *Combinatorics, Probability and Computing*, **30**(2): 200–227, 2021.

[4] A. Nicholas Day and Victor Falgas-Ravry. Maker-breaker percolation games II: escaping to infinity. *Journal of Combinatorial Theory, Series B*, **151**: 482–508, 2021.

[5] Vojtěch Dvořák, Adva Mond, and Victor Souza. The maker-breaker percolation game on the square lattice. To appear in *Journal of Combinatorial Theory, Series B*, 2024.

[6] Vojtěch Dvořák, Adva Mond and Victor Souza. The Maker-Breaker percolation game on a random board. Preprint, arXiv ref: 2402.17547, 2024.

[7] Alfred Lehman. A solution of the Shannon switching game. *Journal of the Society for Industrial and Applied Mathematics* **12**(4): 687–725, 1964.

[8] Miloš Stojaković and Tibor Szabó. Positional games on random graphs. *Random Structures & Algorithms*, **26**(1-2):204–223, 2005.

Project 2: Subcube intersection graphs

Given a d-dimensional hypercube $Q_d := \{0,1\}^d$ and an integer k: $0 \leq k \leq d$, let $\text{Subcube}_k(Q_d)$ denote the collection of all k-dimensional subcubes of Q_d. We construct a graph $G_k(d)$ on $\text{Subcube}_k(Q_d)$ by joining two subcubes from $\text{Subcube}_k(Q_d)$ by an edge if they have a non-empty intersection. The resulting *subcube intersection graph* can be thought of as a hypercube analogue of the celebrated and well-studied Kneser graph.

With motivation coming from social choice theory, Markström and I [2] introduced certain random subcube intersection graph models essentially equivalent to *site percolation* on $G_k(d)$ (i.e. to taking random induced subgraphs of $G_k(d)$). The connection to social choice is as follows: suppose a community is faced with d issues, each of which requires a binary response, 0 or 1. The hypercube Q_d then corresponds to the space of all possible responses. A k-dimensional subcube could then represent the subspace of responses acceptable to a particular member of the community who is flexible on k of the issues but has fixed opinions on the remaining $d-k$ ones. With this interpretation, subcube intersections correspond to possible agreements or acceptable compromises betwen different members of the community.

A number of fundamental questions about the random subcube intersection graph models of [2] remain open, including those of connectivity and component evolution for site percolation on $G_k(d)$. Addressing questions not settled in [2] would be highly suitable for an internship project, and would almost certainly involve investigating the separate, unstudied and intriguing problems of deriving chromatic, isoperimetric and spectral results for $G_k(d)$ — in particular obtaining analogues of a seminal result of Lovász [5] on the Kneser graph and of a classical inequality due to Harper [3] in the hypercube in the 'subcube Kneser' setting.

Further possibilities for the internship research project include extending earlier Ramsey- and Turán-type results of Johnson and Markström [4] on subcube intersection graphs, or more recent work on the closely related hypercube spatial voting model of Day and Johnson [1] and its implications for social choice theory, in particular regarding elections and voting in politically polarised societies.

References

[1] A. Nicholas Day and Robert. J. Johnson Equilibria in a Hypercube Spatial Voting Model. Preprint, arXiv ref:2406.18466, 2024.

[2] Victor Falgas-Ravry and Klas Markström. Random subcube intersection graphs I: cliques and covering. *The Electronic Journal of Combinatorics*, pages P3–43, 2016.

[3] Lawrence H. Harper. Optimal numberings and isoperimetric problems on graphs. *Journal of Combinatorial Theory*, 1(3):385–393, 1966.

[4] J. Robert Johnson and Klas Markström. Turán and Ramsey properties of subcube intersection graphs. *Combinatorics, Probability and Computing*, 22(1):55–70, 2013.

[5] László Lovász. Kneser's conjecture, chromatic number, and homotopy. *Journal of Combinatorial Theory, Series A*, 25(3):319–324, 1978.

Project 3: Bootstrap percolation in random intersection graphs

Bootstrap percolation is concerned with the study of a family cellular automata. It first arose in work of Chalupa and Leath and Reich [1] on magnetic phenomena, and has been widely studied as a toy model for the spread of misinformation or of an infection on a network.

Given a graph G, an infection threshold r and a set of initially infected vertices I_0, a bootstrap percolation process on G unfolds as follows: at each time $t \geq 0$, all vertices of G that are adjacent to at least r vertices of the infected set I_t become infected and are added to I_t to form I_{t+1}. The set of eventually infected vertices in this process is denoted by $I_\infty := \bigcup_{t \geq 0} I_t$.

There are many questions one could ask: given an initially infected set I_0 of a given size, how large/how small could the size of the eventual infection $|I_\infty|$ be? How long does the infection take to reach its final state? If the set of initially infected vertices I_0 is selected at random, what is the likely distribution of I_∞? What if the threshold for infection r is not uniform, or if the infection spreads in a random rather than deterministic fashion?

As part of an internship project, I propose to study bootstrap percolation on random intersection graphs, to provide counterparts in that setting to the work of Janson, Łuczak, Turova and Vallier [3] on the Erdős–Rényi random graph model, and of Whittemore [4] and Sarkar and myself [2] on the Gilbert random geometric graph model.

References

[1] John Chalupa, Paul L Leath, and Gary R Reich. Bootstrap percolation on a Bethe lattice. *Journal of Physics C: Solid State Physics*, 12(1):L31, 1979.

[2] Victor Falgas-Ravry and Amites Sarkar. Bootstrap percolation in random geometric graphs. *Advances in Applied Probability*, pages 1–47, 2021.

[3] Svante Janson, Tomasz Łuczak, Tatyana Turova, and Thomas Vallier. Bootstrap percolation on the random graph $G_{n,p}$. 2012.

[4] Alyssa Whittemore. *Bootstrap percolation on random geometric graphs.* PhD thesis, The University of Nebraska-Lincoln, 2021.

Project 4: Mixed degree conditions for hypergraphs

An r-uniform hypergraph (or *r-graph*) is a pair $H = (V, E)$, where $V = V(H)$ is a set of vertices, and $E = E(H)$ is a collection of r-sets from V that form the edges of H. The *degree* $\deg(S)$ of a set S in H is the number of edges of H that contain S. The minimum of $\deg(S)$ over all subsets $S \subseteq V$ of size $|S| = \ell$ is known as the *minimum ℓ-degree* of H, and is denoted by $\delta_\ell(H)$.

A copy of the clique on t vertices $K_t^{(r)}$ in H is a t-set $X \subseteq V$ such that all $\binom{t}{r}$ possibles r-sets from X are present as edges in H. If H contains no such copy, H is said to be $K_t^{(r)}$*-free*.

A classical problem in extremal hypergraph theory is: given integers $1 \leq \ell < r < t$, determine the maximum value of $\delta_\ell(H)$ over all $K_t^{(r)}$-free r-graphs H on $n = |V(H)|$ vertices. The case $r = 2$ was resolved almost a century ago in a celebrated theorem of Turán [5].

For all $r \geq 3$, however, this problem — known as the *hypergraph Turán problem* —- has remained stubbornly open in all cases, and is arguably one of the major open problems in discrete mathematics; see Keevash's survey [1] dedicated to it.

As part of an internship, I proposed to study a mixed-degree version of the hypergraph Turán problem. The focus would be on obtaining good lower bound constructions for this new family of problems, which is as yet unstudied, in order to cast light on the hypergraph Turán problem. Part of the work would build on a mountain of unpublished data on flag algebraic bounds computed by Sliacan, and on a study of the treasure trove of constructions provided by Sidorenko in his 1995 survey [2].

References

[1] Peter Keevash. Hypergraph Turán problems. *Surveys in combinatorics*, 392:83–140, 2011.

[2] Alexander Sidorenko. What we know and what we do not know about Turán numbers. *Graphs and Combinatorics*, 11(2):179–199, 1995.

[3] Paul Turán. On an extremal problem in graph theory. *Mat. Fiz. Lapok*, 48:436–452, 1941.

Project 5: Covering problems in hypergraphs

An r-uniform hypergraph (or r*-graph*) is a pair $H = (V, E)$, where $V = V(H)$ is a set of vertices, and $E = E(H)$ is a collection of r-sets from V that form the edges of H. The *degree* $\deg(S)$ of a set S in H is the number of edges of H that contain S. The minimum of $\deg(S)$ over all subsets $S \subseteq V$ of size $|S| = \ell$ is known as the *minimum* ℓ*-degree* of H, and is denoted by $\delta_\ell(H)$.

A copy of the clique on t vertices $K_t^{(r)}$ in H is a t-set $X \subseteq V$ such that all $\binom{t}{r}$ possibles r-sets from X are present as edges in H. A $K_t^{(r)}$*-covering* in H is a collection of copies of $K_t^{(r)}$ that together cover all the vertices of H. If these copies may in addition be chosen to be vertex-disjoint (which implies in particular that $|V(H)|$ must be divisible by t), then we say that H contains a $K_t^{(r)}$*-tiling* (also known as a $K_t^{(r)}$*-factor*).

A natural question in extremal hypergraph theory is: given $1 \leq \ell < r < t$, what is the least d such that every r-graph on n vertices with minimum ℓ-degree strictly greater than d must contain a $K_t^{(r)}$-covering? We denote the answer to this question by $c_\ell(n, K_t^{(r)})$, and refer to it as the ℓ*-degree covering threshold* for $K_t^{(r)}$.

The systematic study of covering thresholds was initiated in works of Zhao and myself [2] (in the case $r = 3$, $\ell = 2$) and Markström, Zhao and myself [1] (in the case $r = 3$, $\ell = 1$), and has attracted increasingly significant attention in the last few years. Covering thresholds lie between two of the most studied families of parameters in extremal hypergraph theory, namely the existence thresholds studied in Turán-type problems (so-named after the pioneering work of Turán [5] discussed in the previous subsection) and the tiling thresholds from Hajnal–Szemerédi-type problems (so-named after a major result of Hajnal and Szemerédi [3] from the 1970s), neither of which are satisfactorily understood as of yet — see the surveys of Keevash [4] and Zhao [6] respectively.

Much of the attention has focussed on the 3-uniform ($r = 3$) case, with a scattering of ad hoc results but little yet in terms of an overarching theory. Goals for an internship project would include considering the case of general $r \geq 3$, constructing new examples, and giving general bounds and plausible conjectures for the asymptotic behaviour of the covering thresholds.

References

[1] Victor Falgas-Ravry, Klas Markström, and Yi Zhao. Triangle-degrees in graphs and tetrahedron coverings in 3-graphs. *Combinatorics, Probability and Computing*, 30(2):175–199,

2021.

[2] Victor Falgas-Ravry and Yi Zhao. Codegree thresholds for covering 3-uniform hypergraphs. *SIAM Journal on Discrete Mathematics*, 30(4):1899–1917, 2016.

[3] András Hajnal and Endre Szemerédi. Proof of a conjecture of P. Erdős, in âĂIJcombinatorial theory and its applicationsâĂİ(P. Erdős, A. Rényi, V. Sós, eds), 601-623, 1970.

[4] Peter Keevash. Hypergraph Turán problems. *Surveys in combinatorics*, 392:83–140, 2011.

[5] Paul Turán. On an extremal problem in graph theory. *Mat. Fiz. Lapok*, 48:436–452, 1941.

[6] Yi Zhao. Recent advances on Dirac-type problems for hypergraphs. Pp 145–165 in, *Recent trends in combinatorics*, Springer, 2016.

Gerold Jäger

Gerold Jäger is an Associate Professor of Discrete Mathematics. His research areas are discrete mathematics, combinatorial optimisation, graph theory and algorithmics. It includes the Travelling problem, game theory (especially Mastermind), theory of Latin squares and sensitivity analysis.

Webpage: https://www.umu.se/en/staff/gerold-jager/

Project 1: Travelling salesman problem

This is the problem of finding a shortest route visiting a given set of cities exactly once, called a tour. It is one of the most famous problems in combinatorial optimisation with many applications in transport and logistics. In this project, we develop the theoretical basis and implement heuristical algorithms, i.e., algorithms that find a short, but not necessarily the shortest tour. The used methods lie in stability analysis, i.e., how stable connections are in related problems.

Project 2: Mastermind

Mastermind is a popular board game with two players, the codemaker and the codebreaker. In the original version, the codemaker provides a code consisting of four pegs with six colours for each peg. The codebreaker has to make guesses, and each of them gets an answer from the codemaker on how close (s)he is to the correct code. Then, after a given number of guesses, the codebreaker has to provide the correct code. This project is about investigating and analysing strategies of the codebreaker to minimise the number of guesses. We consider different variants of Mastermind, e.g., when the codebreaker has to give all guesses at once (Static Mastermind) or when the code maker can give partially wrong information (Mastermind with a lie).

Mathematical Modelling and Analysis

Åke Brännström

Åke Brännström is a Professor of Interdisciplinary Mathematics. Most of his work concerns ecological and evolutionary dynamics, and typically adressed questions related to biodiversity and cooperation. He is also interested in mathematical finance.

Webpage: https://www.umu.se/en/staff/ake-brannstrom/

Project 1: Exploring multiplicative payoff functions in the evolution of cooperation

Traditionally, the evolution of cooperation is understood in terms of two distinct strategies: cooperation and defection. Recent advancements, however, recognize this as a spectrum, with full cooperation and defection at opposite ends. Within this spectrum, the possibility of evolution towards purely cooperative or defective behaviors has also been identified.

Most of this research presumes a uniformly mixed population, meaning every individual has an equal chance of encountering any other. Yet, Gylling and Brännström's 2018 study proposed a nuanced approach. They factored in the likelihood of an individual encountering others with similar strategies. Their focus was on additive payoff functions, represented by $B + C$, where B symbolizes the benefit and C denotes the cost.

Building on this foundation, our project aims to delve deeper by examining multiplicative payoff functions, characterized as $B * C$. Past research, particularly by Brännström and colleagues in 2005 and 2011, hints at an intriguing prospect: in certain environments, the tendency to associate with like-minded peers (assortment) might stimulate strategy diversification. This phenomenon has been conspicuously absent in the context of additive payoff functions.

For this project, we will engage in mathematical analyses, implement, and explore individual-based models.

Project 2: Analyzing risk and return of Swedish investment vehicles

Since 2012, Sweden's private investors have been presented with a choice for their stock market investments: the traditional custody account (AD) or the investment savings account (ISK). Each comes with its distinct tax implications:

AD: Investors are taxed 30% on all profits realized within a year, but can potentially recover between 21-30% of their realized losses.

ISK: Investors pay a fixed tax rate derived from the government interest rate, irrespective of whether their assets appreciate or depreciate in value.

At a glance, the ISK might appear to offer superior returns, but it's accompanied by heightened risk. In this project, we aim to dissect these two investment vehicles based on risk and anticipated return. Our goal is to leverage various metrics for risk-adjusted return to discern the situations in which one account might be more advantageous than the other.

Additionally, the AD enables investors to adopt responsive strategies, through which assets can be turned over by selling and buying when market prices fall below the assets' initial purchase values. This approach allows investors to recover part of the realized loss, which can then be channeled into new investments.

This project is supervised by Prof. Åke Brännström, Professor of Interdisciplinary Mathematics, and Prof. Christian Ewald, Professor of Financial Mathematics, both at Umeå University.

Christian Ewald

Christian Ewald is a Professor of Financial Mathematics with specialisation in Financial Mathematics, Quantitative Finance, Risk Management, and Commodities. He is also affiliated with the HINN Business Analytics Group (Norway) as well as the Adam Smith Business School, University of Glasgow (UK).

Webpage: https://www.umu.se/en/staff/christian-ewald/

Project 1: Inflation-indexed securities and derivatives: Applications in fund management

Inflation-indexed securities are securities whose cash flows are linked to the level of inflation. They fulfill three key functions: First, for governments, they reduce borrowing costs as they allow to avoid a costly inflation risk premium. Second, for investment funds, in particular pension funds, they allow for inflation hedging. Third, with a good pricing model, they allow to make relative short-term forecasts of inflation levels, where macro-economic models often fail. In this project prices as well as relevant indices for inflation-indexed securities will be used to make short-term inflation forecasts for different regions. In a continuation of this project or possibly a second project, the impact of short-term inflation on currency exchange rates can be investigated.

Antti Perälä

Antti Perälä is an Associate Professor in Mathematics. His research focuses on functional analysis, complex variables and operator theory. In particular, he is interested in various concrete operators (Toeplitz, Hankel, Volterra, etc.) on spaces of analytic functions. Common topics are their operator-theoretic properties and their spectra, as well as the related optimisation problems. He is also interested in applications to other areas of mathematics and science.

Webpage: https://www.umu.se/en/staff/antti-perala/

Project 1: Hardy, Bergman and related spaces

The study of various spaces of analytic functions is a topic active ongoing research. The most common spaces are different variants of Hardy, Bergman, BMOA and Bloch spaces, as well as the Bargmann-Fock spaces of entire functions. The basic theory of all of these provides good topics for internship projects that can be adapted to the student's interests and ambitions. For instance, basic properties of weighted Bergman spaces would be a good topic blending both complex and functional analysis. The structure of these spaces ensures the existence of the Bergman kernel function, whose properties are central for many problems related to operator theory and optimisation. In many cases, this function is very difficult to calculate in closed form. Therefore, a good project could be to investigate what is known and how the lack of exact information on the Bergman kernel can be overcome.

Project 2: Carleson measures

Carleson measures are named after the Swedish mathematician Lennart Carleson, who used them in his solution to the Corona Problem. A non-negative Borel measure μ is called a p-Carleson measure for an analytic function space X, if X can be continuously embedded into L^p_μ. Characterisation of Carleson measures is a central problem in the study of analytic function spaces, and the solution is known for many different cases. A good internship project would be to study what is known and whether there are some interesting open questions.

Carleson measures have many connections to concrete operators, such as Volterra, Hankel, Toeplitz and composition operators. These topics could very well be incorporated into the project if the student so wishes. The same goes for studying the Corona Problem, or some of its variants.

Project 3: Topological vector spaces

A vector space V with a topology τ is called a topological vector space (TVS) if the vector summation and multiplication by a scalar are continuous with respect to this topology. Particularly important examples of TVS are locally convex vector spaces, which can be seen as generalisations of normed vector spaces. Many of the classical theorems for infinite dimensional Banach spaces, such as the Hahn-Banach Theorem, Uniform Boundedness Principle, and the Open Mapping Theorem, hold true for locally convex vector spaces. Investigating their proofs and other properties of these spaces would be a good internship project.

Many properties of TVS depend on the Axiom of Choice and Zorn's lemma. Another project could be studying TVS from this perspective. A good topic could be to investigate filters and ultrafilters and their connection to the topology in this setting.

Mathematical Statistics

Per Arnqvist

Per Arnqvist is an Associate Professor of Mathematical Statistics. His research areas are functional data analysis and spatial statistics. His research focuses on developing a methodology for model-based clusterization, which includes modelling, prediction and nonparametric inference, with application areas within climate reconstruction and ecology.

Webpage: https://www.umu.se/en/staff/per-arnqvist/

Project 1: Marital fertility modelling

Marital fertility has been studied using age-specific fertility rates, with models such as the Coale-Trussell fertility intensity specification. An alternative approach, named the Belyaev-Waiting model, relies on individual fertility data and introduces the concept of waiting time between pregnancies. This model allows us to better understand marital fertility intensity as the intensity of a married woman becoming pregnant. The project aims to develop bootstrap and resampling methods to estimate the distributions and precision of the parameters within the Belyaev-Waiting model.

Mehdi Moradi

Mehdi Moradi is an Associate Professor of Mathematical Statistics. He specializes in spatial and spatio-temporal point processes, change-point detection, and trajectory analysis. His research focuses on developing advanced mathematical models, that utilize change-point detection, to gain insights into the dynamics of point processes and trajectories.

Webpage: https://www.umu.se/en/staff/mehdi-moradi/

Project 1: Change-point detection for point processes

Significant advancements in data collection and storage capacities have led to vast spatial point data availability from diverse sources, which often give access to (precise) spatial locations and time occurrences of events together with further valuable point-specific information. Applications, among other things, include e.g. locations of traffic accidents and street crimes. The nature of correlations between these events can fluctuate over time, hinting at varying forms of interaction among events during different time intervals. This project is dedicated to the identification of such changes and the classification of point patterns through second-order analysis.

Project 2: Edge correction for Voronoi-based intensity estimators

The spatial distribution of point patterns, which give locations of events such as e.g. traffic accidents and street crimes, is usually estimated via kernel- and Voronoi-based approaches. It has recently been shown that a resample-smoothed version of Voronoi estimators outperforms kernel-based estimators in terms of both bias and variance. However, that estimator suffers from border issues leading to under/over estimations near the borders of the observed area. This project works on proposing an edge-corrected version of the resample-smoothed Voronoi estimators.

Project 3: Multivariate SDE-based models for trajectories

Analysing traffic data is crucial for optimising transportation systems, enhancing safety, and reducing congestion. It provides valuable insights into traffic patterns, allowing authorities to make informed decisions about infrastructure improvements, traffic management strategies, and public transit enhancements. This project focuses on developing multivariate SDE-based models to study how different moving objects interact with each other while taking the geometry of the underlying network into account.

Project 4: Shape-constrained-based estimators for higher-order summary statistics for spatial point processes

Point processes are commonly used to model data that appear as points distributed across a given state space, such as R^2 or linear networks. These models have applications in various fields, including forestry, criminology, and geology, where a key objective is to analyze interactions between points—determining whether they exhibit clustering, inhibition, or random distribution patterns. This project aims to develop shape-constrained estimators for higher-order summary statistics in spatial point processes, addressing data loss issues in existing estimators caused by minus-sampling-based methods.

Project 5: Zero-inflated models for points processes on linear networks

Over the past decade, point processes on linear networks have gotten lots of attention, as in many applications of point processes, events actually happen on some network structures, such as traffic accidents, which can only occur on streets. However, in large networks, many of the segments of the network structure do not host any event, resulting in many zeros. This project focuses on developing zero-inflated models for point processes on linear networks.

Natalya Pya Arnqvist

Natalya Pya Arnqvist is an Associate Professor of Mathematical Statistics. She has two main research interests: statistical regression modelling and functional data analysis. She is specifically interested in developing methods for shape preserving smoothing within generalized additive models and applications of shape constrained additive models (SCAMs).

Webpage: https://www.umu.se/en/staff/natalya-pya/

Project 1: Shape constrained generalized additive models for complex datasets

Whether we're examining how changes in water quality affect fish populations, estimating the lifespan of electronic gadgets, or exploring the link between exercise and heart health, we often rely on statistical modelling to provide us with valuable insights. When trying to analyze how a response variable connects with various predictors using flexible semi-parametric regression models, we might encounter a challenge. Sometimes, these models can become overly flexible, producing results that seem unrealistic. It is often logical to believe that some relationships follow certain shape constraints, such as monotonicity or convexity. The project aims to develop methods within the shape-constrained generalized additive modelling framework that can incorporate large datasets and a wider range of more complex relationships.

Project 2: Functional clustering methods for climate reconstruction

The growing concern regarding global warming and its associated impact on future climate patterns has heightened the need to comprehend historical climate variations over the last centuries and millennia. Understanding past climate variability is crucial for accurate climate modelling and predictions. Yet, there remains a significant knowledge gap, especially concerning changes in seasonal climate, particularly during winters, which recent climate change has disproportionately affected. Varved (annually laminated) lake sediments possess the potential to provide vital insights into past seasonal climates, thanks to their inherent annual time resolution and within-year seasonal patterns. However, the limited availability of statistical tools for quantifying and comparing these seasonal patterns has hindered their full utilisation. This project aims to develop statistical functional clustering methods, particularly tailored to the challenging complexities of varved lake sediment. The implementation of these novel functional clustering techniques, both nonparametric and parametric, will be executed within the R software environment. The project will apply the functional clustering methods to varved sediment data from lakes in Northern Sweden and Finland to infer seasonal climate changes over the last millennia.

Project 3: Advancing cosmic ray analysis: an open-source android neutron detector module

Our planet is continuously immersed in a dynamic flux of cosmic rays, spanning an extensive energy range from 10^7 eV/nucleon to 10^{20} eV/nucleon. Furthermore, the cosmic ray flux exhibits a remarkable variability, fluctuating from merely one particle per square meter and second to a solitary particle per square kilometre and century. The cosmic ray flux at Earth's orbital region is influenced by solar activity, particularly noticeable at energy levels below several tens of GeV/nucleon. This phenomenon positions temporal fluctuations in cosmic ray flux as invaluable indicators of solar activity. Ground-level instruments, such as neutron monitors, are pivotal in observing these fluctuations. Neutron monitors primarily capture secondary cosmic rays, predominantly neutrons, generated when cosmic rays interact with the Earth's atmosphere. This unique interplay between cosmic rays, solar activity, and Earth's environment has urged international collaborations to create global detector networks. The Neutron Monitor Database (NMDB) is a great demonstration of these efforts, with a key objective being global coverage. Achieving this goal entails the establishment of new monitoring stations to bridge the existing gaps. However, installing these stations can present some challenges, making cost-effective alternatives an attractive target. In light of these considerations, this project aims to conceive and rigorously evaluate an open-source Android application module designed to detect neutrons and monitor count rates from a compact neutron monitoring device. The aim is to perform a comprehensive comparative analysis using several statistical approaches, including penalized regression techniques. This analysis will also compare the performance of the proposed neutron detector with some of the existing NMDB neutron monitors.

Patrik Rydén

Patrik Rydén is a Professor of Mathematical Statistics. His research focuses on developing and applying advanced methods for analysing large and complex datasets, with applications in health sciences, life sciences, and the manufacturing industry. His methodological expertise includes supervised and unsupervised machine learning, robust optimisation, and digital twin technologies.

Webpage: **https://www.umu.se/personal/patrik-ryden/**

Project 1: Predicting surgery times using machine learning

In collaboration with the University Hospital of Northern Sweden, this project aims to develop data-driven decision support tools for surgery, post-surgery, and subsequent care pathways. The ultimate goal is to optimise patient throughput and improve the quality of care. The project leverages machine learning to predict surgery durations using preoperative data such as medical records, prescriptions, diagnoses, clinical measures, and planned surgical procedures. These predictive tools are intended to support hospital staff in planning and resource allocation, thereby enhancing both efficiency and patient outcomes.

Project 2: Predicting ambulance driving times using machine learning

The demand for ambulance services is increasing as the number of emergency calls rises with an ageing population. In collaboration with the ambulance service in Northern Sweden, we are developing a digital twin of prehospital care to optimise resource allocation—for example, by adapting staff schedules and strategically locating ambulance stations. The project focuses on applying machine learning techniques to predict ambulance driving times using a range of data sources, including alarm location, ambulance starting point, time of day, weather conditions, predicted routes, and alarm priority. These predictive models will be integrated into the digital twin, enabling data-driven decision support to improve response times and overall efficiency in emergency care.

Project 3: Deriving p-values for predictive models

Supervised classification of large and complex datasets is widely used in many applications. Typically, data are divided into training and test sets, producing a predictive model along with an estimated performance score, such as classification accuracy. We define a model's p-value as the probability of achieving at least the observed performance by chance. The aim of this project is to develop a general approach for estimating such p-values that can be applied across a wide range of classifiers and performance metrics, including black-box algorithms. This work seeks to provide a statistically rigorous framework for assessing the significance of predictive models, thereby strengthening their reliability and interpretability in practical applications.

Jun Yu

Jun Yu is a Professor of Mathematical Statistics with a focus on Statistical Learning and Inference for Spatiotemporal Data, including fundamental research and applications in AI.

Webpage: https://www.umu.se/en/staff/jun-yu/

Project 1: Bayesian hierarchical model for blind source separation of biomedical signals

Your brain plans, organises, and sends commands to your muscles via the spinal cord and the motor neurons. By increasing your force, more motor neurons will be activated, and their activity will increase individually. This activity can be measured from the skin using electrodes by decomposing the electrical activity using so-called blind source separation algorithms. This information can further be used for studying motor control, neuromuscular diagnosis, and rehabilitation purposes. The state-of-the-art blind source separation algorithms within this field are data-driven and exploit the sparse properties of the underlying motor neuron activities. However, current methods do not exploit the rich a priori information about these underlying sources and require manual post-processing to overcome the non-quantifiable estimation uncertainty. This project aims to develop Bayesian hierarchical models to tackle uncertainty estimation problems. The developed models will be compared against the state-of-the-art methods and evaluated using sophisticated simulation models and experimental data with available ground truth data.

Some prior experience with statistical signal analysis, spatiotemporal statistics, and Bayesian hierarchical models is beneficial. So are programming skills in Python, R, or MATLAB.

This is a collaborative project with Dr. Robin Rohlén at the Department of Diagnostics and Intervention: https://www.umu.se/en/staff/robin-rohlen/

Mathematical Foundations of Artificial Intelligence

Axel Flinth

Axel Flinth is an Assistant Professor in Mathematics. His research concerns different kinds of machine learning, i.e. how we with the help of mathematics can help computers find patterns in large sets of data. He is especially interested in compressed sensing, which concerns methods utilizing structrial assumptions to aid reconstruction of signals from incomplete data, and equivariance in deep neural networks, which studies how symmetries in the data can be utilized, and how what effect these methods have.

Webpage: https://www.umu.se/en/staff/axel-flinth/

Project 1: Equivariance and augmentation

In machine learning, the goal is often to detect symmetries 'hidden' in the data. In some tasks, these are however appearant a-priori. As an example, when classifying images of hand-written digits, the classification should not change if the digit is shifted in either x or y-direction. This task is hence *invariant* to translations. If the task on the other hand is to detect *where* the digit is located, the prediction should translate with a translation of the image. Such tasks are *equivariant*.

Under the umbrella of *Geometric Deep Learning*, researchers have come up with methods to, given a group of symmetries, design special types of so called neural networks that are guaranteed to be equivariant. Put simply, the strategy is to restrict the way the linear parts of the networks can be chosen. A more organic approach is to artificially inject the symmetries into the data via so-called *augmentation*. In this project, the aim is to understand the similarities and differences of these two approaches – in particular in the dynamics of the training of the networks. An internship could both involve studying theoretical issues, as well as performing experimental work.

Project 2: Total variation regularized optimization problems on measure spaces

Imagine an object characterized by a few point-like sources of different amplitudes situated somewhere in a continuous domain $\subseteq \mathbb{R}^d$ for some d. One can for instance think of a few bright stars on a night sky, or a collection of small colonies of luminescent bacteria on a petri dish. Such objects can be faithfully modelled by a measure $\mu_0 = \sum_{i=1}^{s} c_i \delta_{x_i}$, where $c_i \in \mathbb{R}$ are the amplitudes, $x_i \in$ are the positions and δ is the Dirac delta distribution. Also imagine that we are given an indirect measurement $b = A\mu_0 + n \in \mathbb{R}^M$ of the object, where A is a linear map and n is a noise vector. Although A surely not is injective (it maps from an infinite dimensional space of measures into the finite-dimensional $\mathbb{R}^M$), μ_0 can still be recovered from b. The key here is that μ_0 has a very special structure in that it only consists of a few spikes. Therefore, μ_0 can be (approximately) recovered from b via the solution of optimization problems regularized by the total variation norm (which is the 'measure version' of the ℓ_1-norm in finite dimension). This project is about analyzing and developing new methods for solving such problems.

Fredrik Ohlsson

Fredrik Ohlsson is an Associate Professor in Mathematics. His research explores the role of symmetries and geometric structures in machine learning. He is particularly interested in the use of differential equations to describe machine learning models and their training.

Webpage: https://www.umu.se/en/staff/fredrik-ohlsson/

Project 1: Stability of neural differential equations

An important problem in machine learning is to quantify and control the stability of models, i.e., the extent to which a small change in the input yields a small change in the output. In neural ordinary differential equations (NODEs), the models are non-linear and high-dimensional dynamical systems, which implies that the dynamics are expected to be non-trivial, involving, e.g., bifurcations and coexisting attractors.

The analysis of the stability of attractors (e.g., equilibrium points, limit cycles, quasiperiodic or strange attractors) in the dynamical systems appearing in NODEs naturally splits into local and nonlocal analysis. The local stability approach, usually based on linearizations, yields information on how the system reacts to arbitrarily small perturbations in a neighbourhood of an attractor. However, a finite perturbation may push a locally stable state into another attractor having a completely different behaviour, which is possibly detrimental for the performance of the NODE. The local stability analysis should therefore be complemented by a nonlocal investigation to quantify the stability of the attractor under investigation.

A classical approach is to construct so called Lyapunov functions in order to characterize the stability of, e.g., equilibrium points and to prove bounds on the size and the shape of the basins of attraction. This project will explore the theory of Lyapunov stability applied to NODEs, and can include both theoretical investigations and numerical experiments.

Alp Yurtsever

Alp Yurtsever is an Assistant Professor in Optimization and Machine Learning. His research focuses on developing theoretical foundations and algorithms for large-scale optimization problems with applications in machine learning and data science. His work addresses challenges such as vast datasets potentially distributed across extensive networks with restricted access, high-dimensional parameter spaces, and complex problem models involving non-smooth and non-convex loss functions.

Webpage: https://www.umu.se/en/staff/alp-yurtsever/

Project 1: Privacy aspects of Frank-Wolfe-type updates in federated learning

Co-supervisor: Ali Dadras

Ali Dadras has recently defended his PhD thesis in mathematical statistics supervised by Alp Yurtsever. His research focuses on personalized models and optimization in federated learning.

Webpage: https://www.umu.se/en/staff/ali-dadras/

The demand for privacy-preserving approaches in Machine Learning (ML) has become increasingly critical in recent years.. Federated Learning (FL) has emerged as a distributed ML paradigm designed to address these privacy concerns while enabling collaborative model training across a large network of clients [1]. By shifting the computational burden to edge devices, FL allows models to be trained locally on private data, making it a vital framework for solving large-scale optimization and learning problems in a distributed manner while upholding privacy principles.

However, FL is associated with significant computational and communication costs. One of the primary objectives in the field is to design optimization algorithms capable of efficiently handling a large number of parameters, particularly in complex learning tasks where parameters are subject to constraints.

In our recent work [2], we introduced a projection-free algorithm called FedFW for constrained FL problems. The FedFW algorithm aggregates Frank-Wolfe (FW) step directions at the server, providing low per-iteration computational cost and facilitating the communication of sparse signals. Preliminary results on the privacy of communicating FW steps instead of gradients are reported in [2], demonstrating the superiority of sharing FW steps over gradient inversion attacks—such as the Deep Leakage algorithm [3]—for the CIFAR100 dataset.

This internship project will focus on a deeper investigation of the privacy aspects of sharing FW steps in an FL setting. The intern will theoretically characterise the privacy benefits of FW-based updates and extend the existing empirical analysis to evaluate their robustness under different adversarial models.

References

[1] Jakub Konečný, H Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.

[2] Ali Dadras, Sourasekhar Banerjee, Karthik Prakhya, and Alp Yurtsever. Federated Frank-Wolfe algorithm. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 58–75. Springer, 2024.

[3] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019.

Project 2: Convex reformulations via copositive programming

Co-supervisor: Karthik Prakhya

Webpage: https://www.umu.se/en/staff/karthik-prakhya/

Karthik Prakhya is a last-year PhD student in mathematical statistics supervised by Alp Yurtsever. His research focuses on machine learning and knowledge discovery in databases.

Copositive programming is an important framework in convex optimization as it allows reformulating many non-convex optimization problems of interest as convex optimization problems in an exact sense [1]. However, despite this convexity, these problems remain intractable due to the complexity of the underlying cone constraint [2]. Recent advancements in quantum annealers offer promising directions to address this intractability by providing new tools for solving certain classes of optimization problems on this cone efficiently [3].

This internship project will explore convex reformulations of selected non-convex optimization problems through copositive programming and evaluate the efficiency of existing solvers. The intern will also investigate alternative relaxations of these formulations that can be efficiently solved using classical computing methods, comparing their performance with quantum-based approaches.

References

[1] Mirjam Dür. Copositive programming–a survey. In *Recent Advances in Optimization and its Applications in Engineering: The 14th Belgian-French-German Conference on Optimization*, pages 3–20. Springer, 2010.

[2] Immanuel M Bomze, Mirjam Dür, Etienne De Klerk, Cornelis Roos, Arie J Quist, and Tamás Terlaky. On copositive programming and standard quadratic optimization problems. *Journal of Global Optimization*, 18:301–320, 2000.

[3] Alp Yurtsever, Tolga Birdal, and Vladislav Golyanik. Q-fw: A hybrid classical-quantum frank-wolfe for quadratic binary optimization. In *European Conference on Computer Vision*, pages 352–369. Springer, 2022.