

Predictive Bayesian inference and foundation models (Project #1)

University of Oslo, Department of Mathematics

PhD programme: University of Oslo, Faculty of Mathematics and Natural Sciences

Starting date: not earlier than 2 January 2027 and preferably as soon as possible thereafter.

Read full announcement and apply!

Description

Foundation models offer unprecedented predictive power, yet their connection to statistical principles of inference remains underdeveloped. This project aims to establish a bridge between modern predictive black-boxes and classical Bayesian theory, guided by the nascent theory of Bayesian predictive inference (also known as predictive Bayes or martingale posteriors). This theory inverts the traditional Bayesian pipeline: instead of starting with a model and a prior, one starts with a (sequence of) predictive distributions and deduces a model and a prior from this sequence. This allows for the use of state-of-the-art predictive machines, such as TabPFN, to conduct Bayesian type posterior inference on parameters of interest. Crucially, the project explores this bridge bidirectionally: either leveraging existing foundation models within martingale posterior frameworks or engineering foundation models to natively satisfy predictive Bayes requirements, including martingale properties and their relaxations.

Project duration: three years.

Project requirements

This page describes requirements specific to Project 1. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree in mathematics, statistics, machine learning, computer science, or another relevant field.
- Documented competence in probability, statistics and machine learning
- Programming experience relevant to scientific computing and data analysis

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process:

- Knowledge of one or more of the following areas Bayesian statistics, large-sample theory, martingales, concentration inequalities, neural networks, or foundation models.

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- Associate Professor Emil Aas Stoltenberg, emilas@math.uio.no

Published Jan. 29, 2024 9:33 PM - Last modified June 25, 2026 1:38 PM

Bridging logic, knowledge representation and learning (Project #2)

University of Oslo, Department of Informatics

PhD programme: University of Oslo, Faculty of Mathematics and Natural Sciences

Starting date: not earlier than 2 January 2027 and preferably as soon as possible thereafter.

Read full announcement and apply!

Description

This is an interdisciplinary project on the boundary between formal logic, reasoning, and knowledge representation on the one hand, and statistical, data driven methods on the other.

As machine learning systems become increasingly “black-box,” the need for high performance models that are provably interpretable has never been greater. This interdisciplinary project merges the rigorous precision of formal logic with the predictive power of statistical learning.

Current data-driven methods often struggle to incorporate domain knowledge or handle structured information. Building on existing methods, such as Logic Regression, we aim to move beyond simple models that predict outcomes, but discover and exploit the underlying “logical rules” of a system. Moreover, where these methods are currently limited to Boolean “on/off” features, this project will explore how to extend them to Knowledge Graphs. We will leverage state-of-the-art methods in computational logic that respect pre-defined domain-specific constraints, and at the same time, sharpen predictive power while being transparent. We will at the same time use uncertainty aware probabilistic techniques to handle inferential and predictive uncertainties. The aim of this research is a new class of “Logic-First” learners that are smaller, faster, and more trustworthy.

Project duration: three years. A fourth year may be considered and it will involve 25 % of other career-promoting work. Other career-promoting work may consist of teaching, supervision, and/or research assistance. This is dependent upon the qualification of the applicant and the current needs of the department.

Project requirements

This page describes requirements specific to Project 2. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

- Master’s degree in computer science, mathematics, statistics, or another relevant field.
- Courses taken at bachelor and/or master level in some of the following: Formal Logic, Logic in Computer Science, Knowledge Representation, Statistical inference, Computational statistics, Bayesian statistics, Optimization, Machine Learning, or similar courses
- Documented programming skills, for example in some of Java, Python, C, C++, C#, R, Julia, etc.

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process:

- Documented competence through courses or master thesis work, in knowledge representation, reasoning, formal logic, or knowledge graphs.

- Documented competence through courses or master thesis work, in statistics, data science, optimization, or machine learning.

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- Professor Martin Giese, martingi@ifi.uio.no

Published Jan. 29, 2024 9:34 PM - Last modified June 25, 2026 1:38 PM

Data attribution for large language models (Project #3)

University of Oslo, Department of Mathematics

PhD programme: University of Oslo, Faculty of Mathematics and Natural Sciences

Starting date: preferably in 2026, by agreement.

Read full announcement and apply!

Description

What parts of the training data are responsible for an LLM's overall behaviour, or influence its answers to specific prompts?

This project addresses the critical explainability challenge of attributing the behaviour of Large Language Models (LLMs) to specific components of their training data. Such attribution is crucial for building trust, diagnosing errors, supporting alignment, and ensuring fairness.

In this PhD project, the candidate will develop new, scalable, and informative data attribution methods that move beyond simple scalar scores. By bridging Explainable AI (XAI), machine learning, and Natural Language Processing (NLP) with statistical rigor, the project will explore novel ways to capture fine-grained attribution effects. This includes disentangling factual contributions from linguistic style, and scaling these methodologies to handle the massive pre-training corpora of modern LLMs.

Possible directions include scalable approximations of attribution based on Shapley values or influence functions, surrogate attribution models, and multi-faceted attribution frameworks. The methods will be implemented and evaluated in relevant LLM settings, such as fine-tuning, learning from examples, and experiments with open-weight language models.

Project duration: three years.

Project requirements

This page describes requirements specific to Project 3. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree or equivalent in computer science (with a specialisation in AI, ML or NLP), mathematics, statistics, or a related field.
- Strong programming skills

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process:

- Experience with Natural Language Processing (NLP) and Large Language Models (LLMs).
- Experience with Explainable AI (XAI) in general, and particularly data attribution methods and/or game-theoretic frameworks like Shapley values.
- Programming skills in Python and deep learning frameworks (e.g., PyTorch, JAX, or TensorFlow).

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- [Professor Ingrid K. Glad glad@math.uio.no](mailto:glad@math.uio.no)

Additional information

This is a joint position with [TRUST – The Norwegian Centre for Trustworthy AI](#). The successful candidate will have access to scientific activities, research networks, and collaboration opportunities across both Integreat and TRUST.

Published Jan. 29, 2024 9:34 PM - Last modified June 26, 2026 9:58 AM

Measuring bias in vision–language models (Project #4)

University of Oslo, Department of Informatics

PhD programme: University of Oslo, Faculty of Mathematics and Natural Sciences

Starting date: preferably in 2026, by agreement.

Read full announcement and apply!

Description

Vision–Language Models (VLMs) have rapidly advanced in their ability to generate rich textual descriptions of visual content. Yet their descriptions are not necessarily neutral: when a model maps an image into language, it may also introduce bias in the form of particular framings, emphasis, omissions, or evaluative language. This PhD project will develop methods and datasets for analysing political and moral bias in VLMs. The project is suitable for applicants interested in both technical model evaluation and broader questions about how AI systems represent social and political realities.

Project duration: three years.

Project requirements

This page describes requirements specific to Project 4. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree in computer science with a specialisation in Natural Language Processing.
- Hands-on implementation experience with large language models, beyond the use of web-based chat interfaces.
- Documented knowledge of core machine learning methods and deep neural networks.
- Documented strong Python programming and data manipulation skills.

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process:

- Experience with bias analysis or working on ethical aspects of AI.
- Experience with VLMs.
- Experience from running large-scale machine learning experiments in HPC environments.
- Experience from annotation or dataset creation.
- Experience with deep learning frameworks like PyTorch.

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- [Professor Erik Velldal, erikve@ifi.uio.no](mailto:erikve@ifi.uio.no)

Additional information

This is a joint position with [TRUST – The Norwegian Centre for Trustworthy AI](#). The successful candidate will have access to scientific activities, research networks, and collaboration opportunities across both Integreat and TRUST.

Published June 25, 2026 11:57 AM - Last modified June 25, 2026 1:39 PM

Models and dynamics of machine reasoning (Project #5)

UiT - The Arctic University of Norway, Department of Physics and Technology

PhD programme: UiT - The Arctic University of Norway, Faculty of Science and Technology

Starting date: preferably in 2026, by agreement.

Read full announcement and apply!

Description

The goal of this project is to advance the understanding of machine reasoning and find new ways to make machine learning models smarter and more trustworthy. The focus will be on the structure and dynamics of the inner or "latent" representation of the data and inspired by approaches such as World Models and JEPA. To discover universal mechanisms, the project will cover different data modalities and combine different types of models, LLMs as well as foundation models. The project will involve both experimental, methodological and theoretical work. Depending on the interests of the PhD student, the project can focus more on theoretical, physics-inspired, approaches such as the study of phase transitions, dynamical systems or be more experimental.

Project duration: three years.

Project requirements

This page describes requirements specific to Project 5. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree in machine learning, computer science, mathematics, physics, statistics or another related field.
- Hands-on implementation experience with large language models, or deep learning models.
- Documented strong programming skills in Python (open-source activities, for example a Github account with open-source projects, will be an asset).

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process: Desired qualifications

- Experience with mathematical modelling and/or good physics background

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

Associate Professor Benjamin Ricaud benjamin.ricaud@uit.no

Additional information

This is a joint position with [TRUST – The Norwegian Centre for Trustworthy AI](#). The successful candidate will have access to scientific activities, research networks, and collaboration opportunities across both Integreat and TRUST.

Published June 25, 2026 11:57 AM - Last modified June 25, 2026 1:39 PM

Multi-agent knowledge bases (Project #6)

UiT - The Arctic University of Norway, Department of Physics and Technology

PhD programme: UiT - The Arctic University of Norway, Faculty of Science and Technology

Starting date: preferably in 2026, by agreement.

Read full announcement and apply!

Description

In the broad area of Agentic AI, this project has the aim to set the foundations for efficient AI agents' access to knowledge. The context windows and Retrieval Augmented Generation (RAG) are too limited for investigating large document corporas and reasoning over them. The project is inspired by recent approaches for building knowledge bases and knowledge graphs for chatbots. The project will involve 1) designing new data structures logically connected as graphs for efficient agent navigation and retrieval 2) methods for building/checking/updating the knowledge bases by multiple agents. This can lead to designing new efficient (shared) memory modules for AI models. The project requires both experimental and theoretical work.

The project is a team work with researchers and engineers of partner institutions who will test the methods on their data.

Project duration: three years.

Project requirements

This page describes requirements specific to Project 6. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree in machine learning, computer science, mathematics, statistics, physics or another related field.
- Hands-on implementation experience with large language models, or deep learning models.
- Demonstrated strong programming skills in Python (open-source activities, for example a Github account with open-source projects, will be an asset).

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- Associate Professor Benjamin Ricaud benjamin.ricaud@uit.no

Additional information

This is a joint position with TRUST – The Norwegian Centre for Trustworthy AI. The successful candidate will have access to scientific activities, research networks, and collaboration opportunities across both Integreat and TRUST.

Published June 25, 2026 11:57 AM - Last modified June 25, 2026 1:39 PM

Probabilistic representation learning (Project #7)

UiT - The Arctic University of Norway, Department of Physics and Technology

PhD programme: UiT - The Arctic University of Norway, Faculty of Science and Technology

Starting date: not earlier than 2 January 2027 and preferably as soon as possible thereafter.

Read full announcement and apply!

Description

Self-supervised learning has become a key technology in modern AI, enabling models to learn powerful representations from large amounts of unlabeled data. Despite their success, many current approaches rely on heuristic techniques for data augmentation and preventing representation collapse, with limited theoretical understanding.

This PhD project aims to develop a principled probabilistic framework for self-supervised learning. The goal is to build more robust, stable, and trustworthy AI models with improved generalization and uncertainty estimation, while providing stronger theoretical foundations for future self-supervised learning methods.

Project duration: Four years. The nominal length of the PhD programme is three years. The fourth year is distributed as 25 % each year and will consist of teaching and other duties.

Project requirements

This page describes requirements specific to Project 7. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree in mathematics, statistics, computer science, data science, machine learning, or a another related field.
- Documented background in statistics, machine learning, and scientific programming (e.g., Python/PyTorch/Tensorflow).

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process:

- Experience with representation learning, probabilistic modeling, and Bayesian methods will be considered an advantage.

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- Professor Michael Kampffmeyer michael.c.kampffmeyer@uit.no

Structured VLMs: panoptic scene graphs for high-level reasoning (Project #8)

University of Oslo, Department of Informatics

PhD programme: University of Oslo, Faculty of Mathematics and Natural Sciences

Starting date: not earlier than 2 January 2027 and preferably as soon as possible thereafter.

Read full announcement and apply!

Description

While current Vision-Language Models (VLMs) have demonstrated impressive capabilities, their reliance on large, unstructured datasets often leads to a superficial understanding of complex visual scenes. This project introduces a paradigm shift by proposing a structured, graph-based approach to vision-language fusion. The central innovation is the creation of panoptic scene graphs, a rich representation that explicitly models the relationships between objects and concepts in an image. These graphs will be derived from a novel application of superpixels and fused with the advanced reasoning capabilities of Large Language Models (LLMs). This bidirectional flow of information enables the system to move beyond superficial pattern matching and engage in genuine compositional reasoning.

The project aims to produce ground-breaking results on tasks requiring deep contextual understanding, develop a more accurate and explainable machine learning methodology, and unlock a new era where AI reasons over images with the depth of a knowledge graph.

Project duration: three years.

Project requirements

This page describes requirements specific to Project 8. Applicants must also satisfy the general qualification requirements, PhD programme admission requirements, and application requirements described in the main call for PhD fellowships.

Required qualifications

Applicants should satisfy the following qualifications to be eligible for consideration for this project:

- Master's degree in mathematics, statistics, computer science, data science, machine learning, or another related field.
- Documented background in mathematics, statistics, machine learning, and scientific programming (e.g., Python/PyTorch/Tensorflow).
- Expertise in creating or adapting modern, token-based transformer architectures.
- Expertise in creating or adapting Large Language Models (LLMs) and interfacing them with other models.

Preferred qualifications

The following qualifications are not required but will be considered an advantage in the evaluation process:

- Familiarity with Panoptic Scene Graph Generation (PSG) methodologies.
- Familiarity with Large Language Model fine tuning and alignment.
- Knowledge of graph-based Retrieval-Augmented Generation (Graph RAG) workflows.
- Understanding of probabilistic modeling and information-based regularizers.

Contact for scientific inquiries

For questions regarding the scientific content, research objectives, or project-specific requirements

- [Professor Adín Ramírez Rivera adinr@uio.no](mailto:adinr@uio.no)

Published Jan. 29, 2024 9:35 PM - Last modified June 25, 2026 1:40 PM