

PhD Retreat

13th–15th October 2026

PROGRAM DRAFT • including talk titles and abstracts

Rooms and restaurant	Hótel Stracta – stractahotels.is
Address	Rangárflatar 4, 850 Hella, Iceland – GPS location
Organizers	Morris Riedel and Gabriele Cavallaro

Day 1 – Monday 12.10.2026

Arrival day for international participants

Bus options from Keflavik airport for international arrivals

Reykjavik Excursions: re.is/tour/flybus

The bus arrives at the BSÍ bus terminal in Reykjavík: [map](#)

The transfer takes about 45 minutes and costs 3999 ISK.

The bus company also offers transfers to/from the airport via the Blue Lagoon: [Blue Lagoon transfer](#)

17:00–19:00

The 27th IHPC workshop

Fenjamýri meeting room in Gróska, Kristínargata 1, Reykjavík

ihpc.is/events/27th-icelandic-hpc-community-workshop

Day 2 – Tuesday 13.10.2026

09:30–11:00	Transport to Stracta Hotel, Hella From Reykjavik center (Hallgrímskirkja) by bus and selected cars
11:00–11:30	Arrival at Stracta Hotel Luggage goes to storage until check-in time
11:30–12:30	Welcome lunch at Stracta Hotel
12:30–12:45	Welcome & Agenda Morris Riedel & Gabriele Cavallaro
PhD student / post doc / research assistant presentations (12:45–17:30)	
12:45–13:25 40 min	Surbhi Sharma <i>PhD defense rehearsal</i> Advancing Land-Cover and Land-Use Classification with Semantically Consistent and Data-Efficient Deep Learning on Multi-Source Earth Observation Data ABSTRACT The growing availability of high-resolution, open-access Earth Observation (EO) data is revolutionizing our capacity to monitor the surface of our planet and its intrinsic properties. This has led to the accelerated development of advanced Machine Learning (ML) and Deep Learning (DL) models tailored for applications in Remote Sensing (RS), such as land-use and land-cover mapping. However, the success of these models is often limited by the need for extensive annotations from a diverse distribution of data and the significant semantic differences that arise between data from various sources and modalities. This PhD study aims to leverage publicly available, large-scale, multi-source datasets, especially satellite RS data and geo-tagged images, to enhance the performance of land-use-land-cover classification systems. These methods are being developed to be scalable on modern High-Performance computing systems, incorporating semantically consistent and data-efficient DL strategies.
13:25–13:35	<i>Coffee break</i>
13:35–13:55 20 min	Liang Tian Assessing the Influence of Spatial and Thematic Resolution of Satellite-derived Images on Landscape Metrics ABSTRACT Evaluating landscape complexity is fundamental for broad-scale ecological and environmental monitoring. While higher spatial resolution is widely assumed to provide better characterizations of terrestrial ecosystems, the parallel role of thematic resolution has been largely overlooked. In this talk, I present our comparative framework across 27 US-based NEON sites using Sentinel-2 (10 m), Landsat 8 (30 m), and MODIS (250 m/500 m) data. We systematically decouple spatial and thematic resolution by analyzing structural image texture alongside categorical landscape metrics, validating their explanatory power against surveyed species richness of breeding birds (750 m) and small mammals (90 m). We identify a clear spatial-thematic trade-off: high spatial resolution is critical for resolving fine-scale physical vegetation structure, but fine thematic resolution is far more decisive for compositional landscape metrics, where 30 m NLCD outperforms the thematically generalized 10 m WorldCover. The talk concludes with methodological takeaways on scale-matching and multi-resolution data selection in Earth observation.
13:55–14:15 20 min	Kennedy Adriko Uncertainty-Weighted Downscaling of District Crop Yield Statistics to Parcels Using Geo-Foundation Model Embeddings ABSTRACT Although crop-yield statistics are widely available at the district level, they obscure significant variability within districts. Parcel-level yield observations provide more detailed spatial information but are limited, geographically concentrated, and often proprietary. Consequently, inferring parcel-level yields from district statistics represents an underdetermined inverse problem, particularly under the constraint that disaggregated predictions must reconcile with published aggregate totals. This study introduces an uncertainty-aware downscaling framework that integrates district-level crop-yield statistics, limited parcel-level observations, and geospatial foundation model embeddings to estimate winter-wheat yields for agricultural parcels in Germany.

Frozen TerraMind embeddings from Sentinel-2 imagery generate parcel-level representations for yield and uncertainty prediction. Sparse harvester observations supervise parcel-level predictions using a scale-adaptive Huber objective. An area-weighted observation operator links parcel estimates to reported district yields. An uncertainty-weighted reconciliation mechanism distributes discrepancies between predicted and reported district yields across parcels based on their estimated uncertainty.

14:15–14:35

20 min

Stefano Maurogiovanni

Timely Crop Phenology Prediction from Field-Level Imagery and Satellite Time Series with Vision–Language Models

ABSTRACT

Accurate and timely reference data are essential for training and assessing satellite-based machine learning models for agricultural mapping and monitoring. Because agricultural systems are highly dynamic, reference data must be continuously updated. At the same time, the expanding scope of applications requires reference data that goes beyond crop type labels. To address these limitations, we propose a novel approach that leverages zero-shot Vision Language Model (VLM) capabilities to automatically extract timely agricultural reference data from geo-tagged field photos, extending coverage beyond manual surveys while minimizing accuracy loss.

14:35–14:45

Coffee break

14:45–15:05

20 min

Sami Hashim

Amortized self-supervised pretraining for compute-efficient environmental monitoring: How well do frozen 3D-structure foundation models transfer?

ABSTRACT

Efficiency claims in Earth Observation (EO) deep learning are usually estimated from floating-point counts rather than measured. We report a metered account of the energy and accuracy of foundation-model adaptation on a Grace-Hopper supercomputer, covering two regression tasks with contrasting input profiles: 10,000 multi-temporal optical and radar chips requiring 25 raster opens per sample, and 22,000 in-situ soil-carbon layers sampled from very large covariate rasters. Each is tested in three adaptation regimes: frozen encoder, full fine-tuning, and a randomly initialised trunk of identical architecture, with per-rank energy tracking separating GPU, CPU and RAM consumption. Accuracy is governed by architecture and adaptation regime; energy is governed by data movement and by whether a backward pass through the trunk is paid for. The two are therefore not co-monotone, and we report the resulting accuracy-versus-energy frontier.

15:05–15:25

20 min

Sayan Mandal

Title: *TBD*

ABSTRACT

15:25–15:45

20 min

Rui Wang

Towards large-scale simulation of the hybrid diffusion model

ABSTRACT

This presentation explores why scaling hybrid quantum–classical diffusion models is necessary to understand whether quantum components can provide meaningful benefits beyond small proof-of-concept settings. We investigate how quantum capacity can be increased while keeping simulation and training computationally tractable, combining architectural design with scalable simulation strategies. Current experiments identify clear efficiency–capacity trade-offs and show that the placement and structure of quantum components strongly affect computational cost. Building on these results, we outline a path toward larger-qubit hybrid diffusion models using more efficient quantum architectures, distributed simulation, and systematic scaling studies to determine when increased quantum capacity translates into improved generative performance.

15:45–15:55

Coffee break

15:55–16:15

20 min

Hao Geng

Iterated Simulated Annealing for Combinatorial Optimization

ABSTRACT

Simulated annealing (SA) typically explores the solution space through a single cooling trajectory. Iterated simulated annealing (ISA) treats a finite SA run as a repeatable search step. Each cycle cools until a freeze criterion is met and then applies a local search to the terminal state. The refined state starts the next anneal, while a passive archive keeps the best solution found. We evaluated ISA on Max-Cut, including the Gset and K2000 benchmarks, and on fixed-K community detection. On the G-

set benchmark, ISA reached the best-known cut on 54 of 71 instances, and it also reached the best-known cut on the fully connected K2000 instance. For community detection, ISA reproduced the maximum modularity reported with Fujitsu's Digital Annealer on Zachary's Karate Club and on the IEEE 33-bus and 118-bus systems. Compared with a single SA run using the same number of sweeps and the same starting and final temperatures, ISA was more reliable at most budget points. On IEEE 118-bus system at the largest budget, 94% of its runs reached the target, against 25% for the single SA run.

16:15–16:35

20 min

Daniele Franch

From MAP to MPM: Quantum Annealing for Sampling-Based Inference in CRFs for Semantic Segmentation with Sparse Ground Truth

ABSTRACT

Conditional random fields (CRFs) are commonly used to refine convolutional neural network predictions in the semantic segmentation of remote sensing images. When the network is trained with sparse ground truth, however, the CRF is only an approximate model of the true labeling, and its minimum-energy (MAP) solution tends to over-smooth the map. In this work, we replace MAP inference with the maximizer of the posterior marginals (MPM), which assigns each pixel the class of highest marginal probability estimated from samples of the CRF posterior, and we draw these samples with a quantum annealer, which naturally acts as a sampler of low-energy configurations. To fit current hardware limitations, the problem is reduced by fixing the pixels whose label is certain and restricting the remaining ones to their two most likely classes, and is then sampled block by block, each block conditioned on the current labels of the rest of the image. Experiments on the ISPRS Vaihingen dataset show that classical MPM inference on the full multiclass problem improves accuracy over the graph-cut MAP solution, and that annealer-based MPM, despite the approximations introduced, also slightly outperforms it. Moreover, all MPM variants, classical and quantum, require less computation time than the graph-cut baseline. These results indicate that, for CRFs trained with sparse ground truth, sampling-based inference is a valid alternative to energy minimization, and one for which quantum annealers are naturally suited.

16:35–16:45

Coffee break

16:45–17:15

Day 2 Closing & Round Table Discussion – Postdoc Experiences

Rocco Sedona, Marcel Aach, Ehsan Zandi, Amer Delilbasic, Vadim Vybornov, Chadi Barakat

17:15–20:00

Free time

20:00–21:30

Dinner at Hotel Stracta

PhD student presentations (09:00–14:15)

09:00–09:30

30 min

Pinaki Banerjee

From Natural Language to Crop Maps: An Agentic Approach to Earth Observation

ABSTRACT

Earth Observation (EO) applications require integration of heterogeneous data sources, computational workflows and foundation or application-specific models, each with distinct input formats and preprocessing requirements. These components are commonly integrated through manually defined pipelines that encode dependencies and assumptions. Adapting such pipelines to other datasets, models, or user requirements can require substantial implementation effort, limiting their flexibility and reuse.

Large Language Models (LLMs) capable of tool use offer a mechanism for translating natural-language objectives into executable workflows. The Model Context Protocol (MCP) [1] provides a standardized interface for accessing external tools and data resources, enabling the separation of workflow orchestration from domain-specific computation. Within this architecture, an LLM selects and parameterizes available tools, interprets execution feedback, and adjusts subsequent operations.

We propose an MCP-based agentic architecture for EO analysis and present an initial prototype supporting workflows built around the TerraMind geospatial foundation model [2]. The system combines a locally running Hermes agent harness [3] with language-model inference accessed through Blablador [4]. Given a natural-language request and a spatiotemporal extent, the agent decomposes the analytical objective into operations executed through a predefined set of MCP tools. Geospatial outputs are computed from EO data using the exposed processing functions and models, while the LLM coordinates execution and result communication.

09:30–09:50

20 min

Vadim Vybornov

Harnessing the Agent: Layered Verification for Autonomous Workflows

ABSTRACT

AI agents, as systems based on LLMs, are probabilistic systems. They can hallucinate, select the wrong tools, report incorrect results, or stop before finishing. Agents' capabilities are known to degrade significantly as their context windows fill. One of the patterns used to avoid context-window degradation during a long-running task is the so-called RALPH loop. The initial task input is split into a queue of self-contained steps executed one at a time, each in a fresh session with a clean context. Each step reads its specific task, performs one bounded unit of work, writes down its outputs, and exits. This approach can provide an operational framework for data pipelines: analyzing pipeline instructions, planning and implementing changes, and triggering real runs.

Fresh-context workers are, however, not reliable on their own. They can stop operating midway, touch files outside the task's scope, or assert facts contradicted by the code they inspected. At the same time, their reports can remain coherent and properly formatted. Neither plain code nor another model can catch them alone: code lacks semantic awareness, while an unconstrained model reviewer often adopts the worker's flawed framing.

To address this, we prototype a layered verification harness built around the RALPH loop, separating execution from acceptance and evidence from authorization. Each operational unit consists of a worker and a judge, each followed by its own deterministic gate — pure code without LLM reasoning. The first gate checks mechanical facts, including process exit, report integrity, protected files, and observed run state. A read-only judge from a different model family assesses the report against a frozen workspace. It is instructed to state task-specific criteria before reading the answer and to record evidence for the claims it checks. The second gate checks the judge's record for contradictions and determines whether to accept the attempt. Rejected attempts are rolled back; feedback and previously refuted claims carry into the next attempt. Unresolved cases go to a human.

The harness does not prove that an agent's output is correct. It provides a practical way to distinguish what code can check, what a model must assess, and what still requires human attention.

09:50–10:00

Coffee break

10:00–10:20

20 min

Anis Abdollahi-Sissan

Learning Compact Representations from Unlabeled ICU Time Series

ABSTRACT

Continuous bedside physiological monitoring in neonatal and pediatric intensive care units generates large amounts of multivariate time series. These recordings could support data-driven clinical decisions, for example by helping to detect clinically relevant events, complications, or changes in

patient state. However, training such data-driven models often requires clinical annotations of relevant events, which are time-consuming to obtain and are usually limited to specific clinical questions or event types. Unsupervised representation learning and clustering are therefore attractive in this setting, because they can exploit large amounts of unlabeled time series while using scarce annotations for evaluation and interpretation. At the same time, these recordings are heterogeneous, high-dimensional, and often sampled at different temporal resolutions across sensors, so raw time series are difficult to compare directly, and clustering them quickly becomes computationally challenging. Together, sparse annotations and computationally demanding raw time series data create two challenges: how to construct compact representations that preserve clinically relevant temporal structure, and how to evaluate such representations when labels are sparse and internal clustering scores do not necessarily reflect clinical usefulness. My research investigates how raw physiological time series can be transformed into compact and informative representations that support clustering, complication detection, and clinical interpretation. In this talk, I will give an overview of three methods for learning or constructing such representations: embeddings learned with encoder models using contrastive objectives, feature-based representations extracted from time-series windows, and shapelet-based methods that capture local temporal patterns. I will also present some of our preliminary findings on embedding-based and feature-based representations, assessed via clustering structure and clinical annotations, and compare the computational cost and interpretability of all three approaches.

10:20–10:40

20 min

Camelia Oprea

AI-based complication detection in neonatal intensive care data

ABSTRACT

Neonatal intensive care units generate large volumes of high-frequency physiological time-series data, including ECGs, airway flow, oxygen saturation, and other continuously monitored signals. These data offer substantial potential for AI-based complication detection, yet neonates, especially preterm infants, are highly vulnerable, physiologically dynamic, and developmentally immature, limiting the direct applicability of many algorithms designed for adults.

In our work, we investigate AI methods for detecting four clinically relevant neonatal complications, including patient ventilator asynchrony, apnea bradycardia syndrome, and continuous monitoring of blood glucose abnormalities. The studied tasks include both univariate and multivariate detection problems, with event durations ranging from less than one second to several minutes or hours. We compare approaches ranging from rule-based methods and random forests to neural networks and foundation models.

Our research focuses on how preprocessing steps, such as signal cleaning and annotation consistency, influence algorithmic performance. We further assess which methods are most suitable for specific complication types, their computational efficiency for continuous parallel bedside monitoring, and their explainability, which can be essential for clinical acceptance and regulatory translation. Finally, we explore how raw detections can be transformed into clinically meaningful scores, supporting the integration of AI-based monitoring into neonatal clinical decision support systems.

10:40–10:50

Coffee break

10:50–11:10

20 min

Syed Muhammad Azhar

Title: *TBD*

ABSTRACT

11:10–11:30

20 min

Zahid Ali

Title: *TBD*

ABSTRACT

11:30–12:55

Lunch break

12:55–13:15

20 min

Greta Bellagamba

A Finite Element Glacial Isostatic Adjustment Model of Iceland

ABSTRACT

Glacial Isostatic Adjustment (GIA) is the deformation the Earth undergoes in reaction to glacier changes. The earth deformation rate is controlled not only by the ice history, but also by the rheology of the Earth, which responds both elastically and viscoelastically. In my PhD project, I use the Finite Element (FE) software COMSOL Multiphysics to create an Earth model which I load with an ice history based on campaign data and Digital Elevation Models, spanning the time period 1890–2023,

considering a steady-state prior 1890. The resulting GIA model predictions of earth deformations are then compared to geodetic data to determine the two best-fit Earth parameters: the viscosity of the mantle and the lithospheric thickness. The GIA-induced pressure changes in the mantle lead to magma generation through a process called decompression melting. We compute the melt produced by this phenomenon from 1890 to 2023, obtaining a total produced volume equal to 27 km³, the equivalent of two Laki Eruptions (1783–1784). Future projections of glacier development based on different climate scenarios will be used to predict future earth deformation rates and help in infrastructure risk assessment and planning.

13:15–13:55

40 min

Joint presentation – Ehsan Zandi & Edoardo Pasetto

Ehsan Zandi

Title: *TBD*

ABSTRACT

Edoardo Pasetto

ISAC: Fundamentals and Applications in Multimodal Environment Reconstruction

ABSTRACT

Integrated Sensing and Communication (ISAC) is emerging as a key technology for future wireless networks, enabling the joint use of communication and sensing capabilities to enhance system performance over a variety of applicative tasks.

This presentation is divided into two parts. The first part provides a general introduction to ISAC, covering its fundamental concepts, key principles, and potential applications. The second part focuses on multimodal environment reconstruction as a practical use case. In particular, it explores how information from multiple sensing and data sources can be combined to construct accurate 3D representations of the surrounding environment. These reconstructed models can subsequently be used for ray-tracing based radio propagation simulations, supporting more realistic and environment-aware wireless network modeling.

13:55–14:15

20 min

Ingólfur Hjörleifsson

Title: *TBD*

ABSTRACT

14:45–14:55

Meeting in hotel lobby

Preparation for departure to the LAVA Centre

14:55

Transfer to the LAVA Centre by minibus

Group divided in 2 for transfer

15:15

LAVA Centre visit

17:00

Transfer back to Stracta

17:35–20:00

Free time

20:00–21:30

Dinner at Hotel Stracta

Day 4 – Thursday 15.10.2026

07:30–09:00	Breakfast at Hotel Stracta restaurant
09:00–10:00	Check-out
10:00–10:30 30 min	Claudia Paris Integrating multimodal satellite data, field photos, and AI to go beyond traditional mapping ABSTRACT Increasingly frequent extreme weather events and climate change pose growing threats to human well-being and ecosystem health. To monitor and assess environmental change and land degradation (e.g., deforestation, desertification, urban expansion), satellite data can be leveraged for the regular production of land-cover maps that capture dynamics across the Earth's surface. These data have been widely used to generate consistent, large-scale land-cover information over time, enabling monitoring at regional to global scales. With rapid advancements in Earth Observation technologies and computational capabilities, there is now an opportunity to move beyond traditional land-cover mapping toward more comprehensive and actionable characterizations of landscape dynamics. In particular, the increasing availability of crowdsourced, geo-tagged field photos introduces an additional and complementary data source. When integrated with satellite imagery, these data enable richer, multi-dimensional representations of land systems. This integration allows information to be captured at multiple levels of semantic detail, which is essential for understanding complex and heterogeneous landscapes. While satellite image time series provide valuable temporal information on land-cover dynamics, ground-level photos offer fine-grained spatial detail and contextual insights into local conditions, such as vegetation structure, land management practices, disturbance events, and small-scale landscape features that are difficult to observe from space. Recent advances in Artificial Intelligence (AI), particularly in visual-language models, self-supervised learning, and multimodal approaches, create new opportunities to effectively combine satellite and ground-based data. The integration of geo-tagged field photos, satellite image time series, and AI methods enables more detailed and scalable interpretation of land-cover patterns and processes, supporting improved detection and understanding of environmental changes that were previously difficult to capture.
10:30–12:00	Closing – Morris Riedel & Gabriele Cavallaro Summary of PhD Retreat objectives and challenge areas, plan for executive summary and next retreat
12:00–13:30	Lunch at Hotel Stracta
14:00–15:30	Transfer back to Reykjavik center
15:30	End of program