Machine Learning Day

3^{th} Tuesday, June

Amphipôle - Auditorium C

Detailed program

09:15-09:30
Lunch Poster

Jaume Banus Cobo, FBM,CHUV, Poster 11

Cardiovascular diseases remain the leading cause of death worldwide. Cardiac imaging is essential for accurately assessing heart structure and function, playing a key role in diagnosis, treatment planning, and prognosis. In this study, we introduce a new way to represent spatio-temporal cardiac data using a multiplex graph and a multi-level message passing neural network to classify patients with different cardiovascular conditions. Using open data from the Automated Cardiac Diagnosis Challenge (N=150), our results demonstrate that this multiplex approach effectively captures important features in the data, allowing us to predict and simulate the spatio-temporal trajectories of cardiac features, such as thickness and volume during the cardiac cycle and for different clinical groups.

Graph Neural Network | Images

09:30-09:45
Lunch Poster

Amro Abdrabo, HEC, Poster 2

The usage of anomaly detection is of critical importance to numerous domains, including structural health monitoring (SHM). In this study, we examine an online setting for damage detection in the Z24 bridge. We evaluate and compare the performance of the elliptic envelope, incremental one-class support vector classification, local outlier factor, half-space trees, and entropy-guided envelopes. Our findings demonstrate that XGBoost exhibits enhanced performance in identifying a limited set of significant features. Additionally, we present a novel approach to manage drift through the application of entropy measures to structural state instances. The study is the first to assess the applicability of one-class classification for anomaly detection on the short-term structural health data of the Z24 bridge.

Decision Trees | Chronological series XGBoost

09:45-10:00
Lunch Poster

Luca Stickley, FBM,CHUV, Poster 13

This study aims to predict lung function trajectories up to 12 months post-severe SARS-CoV-2 using oropharyngeal (OP) and rectal (REC) microbiota data. A graph-embedded deep feedforward network (GEDFN) with multi-task learning (MTL) was developed and applied to the microbiome to predict lung function parameters. The model demonstrated the highest predictive accuracy as measured by the area under the curve (AUC) for FEV1/FVC (>0.99), and strong general prediction (>0.80) for other lung functions measures. Excluding either OP or REC microbiota data significantly reduced the AUCs for FEV1 (>0.8 to <0.6) and FVC (>0.8 to <0.7). Removing MTL decreased AUCs for DLCO and TLC (>0.8 to <0.6). AUCs remained stable across timepoints, with improvements at 12 months for DLCO (0.87) and FVC (0.89). Strong associations between microbiota composition and lung function measures persist up to 12 months post-hospitalization in this cohort, highlighting its potential as a predictive biomarker for long-term pulmonary recovery.

Direct action neural networks | Grpah neural networks | Tabulars | Graphs

10:00-10:15
Lunch Poster

Vladyslav Zalevskyi, FBM,CHUV, Poster 15

Abnormal neurodevelopment affects about one in ten children, often leading to neurological disabilities and significant socioeconomic costs. Although two-thirds of cases could be detected before birth, many are diagnosed late, reducing early intervention opportunities.MRI is increasingly used for fetal brain imaging due to its superior resolution and contrast, offering insights into neurodevelopmental trajectories. However, pediatric neuroimaging faces challenges from domain shifts caused by varying imaging conditions and limited annotated clinical data. Generative models and synthetic data address these issues by simulating diverse imaging scenarios and enhancing model performance despite data limitations. Our project focuses on developing advanced generative techniques to improve fetal neonatal MRI image processing. In it we show how they help overcome domain shifts and data scarcity, enhancing early diagnosis and intervention and making neuroimaging techniques more accessible for both clinical and research applications.

Convolutional neural networks | Generative models | Images

10:15-10:30

Ebi Antony George, FBM/CHUV, No poster

Deep learning has opened up opportunities to quantify animal behaviour at a higher resolution and gain novel insights into the decisions of individuals and collectives. In this talk I will give an overview of several tools that are being used in the emerging field of quantitative animal behaviour through two studies – one laboratory-based and another field-based. In the first study on courtship in Drosophila, I will discuss how deep learning can be used to quantify behaviour from video data using a combination of pose estimation and a transformer-based classifier trained on the video data. In the second study on primate behaviour in the field, I will show work comparing different machine and deep learning algorithms to classify behaviour from accelerometer data obtained from collars attached to vervet monkeys. I will end with a brief discussion on the hurdles that researchers face in implementing these tools in animal behaviour studies.

Random forests | XGBoost | Convolutional nerual networks | Transformers | Tabulars | Videos | Chronological series

11:00-11:15
Lunch Poster

Nataliia Molchanova, FBM,CHUV, Poster 17

"Trustworthy AI for Enhanced Multiple Sclerosis Patient Care" : *Artificial intelligence (AI) is successfully used to automate medical image annotation, a tedious and time-consuming task often subject to high inter-rater variability. This automation is particularly crucial for multiple sclerosis (MS), a chronic neurodegenerative disease. In MS, temporal and spatial dissemination of brain focal lesions is essential for diagnosis, prognosis, and treatment planning. The deep-learning (DL) approaches show the highest accuracy among other methods for automatic lesion segmentation on magnetic resonance images. The black-box nature of DL models and opaque decisions result in a lack of trust by stakeholders and hinder their clinical adoption. Trustworthy AI aims to bridge this trust gap by assessing the reliability of model predictions. Uncertainty quantification, rooted in Bayesian theory, plays a significant role in this assessment. We propose a novel trustworthy AI framework focused on focal lesion segmentation, evaluating model reliability at multiple anatomical scales : voxel, lesion, and brain. In a cohort of 117 patients, we derive an interpretability analysis by leveraging lesion-scale uncertainty values. The study reveals similarities between factors influencing annotator and model confidence. These insights can inform strategies for improving model performance and fostering trust in AI-assisted MS diagnosis and treatment planning.*

Convolutional neural network | Images | 3D magnetic resonance images

11:15-11:30
Lunch Poster

Ilia Azizi, HEC, Poster 6

Accurate prediction of tropical cyclone (TC) intensity is crucial for early warning systems and disaster preparedness but remains challenging due to the complex interactions influencing cyclone development. Traditional machine learning models often rely on single-modal data sources—such as satellite imagery or atmospheric measurements—which may not capture the full spectrum of factors affecting TCs. We present MultiCyclone, a multi-modal machine learning approach that integrates tabular atmospheric data from the Statistical Hurricane Intensity Prediction Scheme (SHIPS) dataset, satellite images processed from European Centre for Medium-Range Weather Forecasts (ECMWF) reanalysis data, and textual weather reports from the National Hurricane Center (NHC) Tropical Weather Discussion (TWDAT) archives. By leveraging the strengths of each data modality, our model enhances predictive accuracy, discovers novel cyclone intensity predictors, and remains computationally efficient for real-time forecasting applications. This project exemplifies the application of advanced ML techniques to a critical real-world problem, offering practical tools for disaster preparedness and mitigation effort.

Linear/logistic regression | Random forests | Convolutional neural networks | Réseaux de neurones récurrents | Transformers Multi-modal neural networks Tabulars | Images | Texts | Chronological series

11:30-11:45
Lunch Poster

Stéphanie Blanc, FBM,CHUV, Poster 18

This study aims to assess the performance of open-source, locally run Automatic Speech Recognition (ASR) for verbatim transcription and Large Language Models (LLMs) for text data analysis of research interviews. Audio-recorded semi-structured interviews (N=12) were conducted in French to evaluate participants' opinions on the usefulness and appropriateness of a web-based application for cannabis harm reduction. The transcription and analysis techniques are implemented using the psifx package, which provides streamlined workflows for integrating ASR systems (OpenAI Whisper) and LLMs (DeepSeek R1, LLaMA 3.1 and 3.3). ASR verbatim transcriptions will be compared with human transcriptions using generation time, diarization errors, and word error rate. LLMs content analyses will be compared with deductive qualitative analyses using evaluation metrics for classification tasks (e.g., accuracy, precision, recall) and extraction tasks (e.g., instruction adherence, faithfulness, relevance). As such, this study assesses the feasibility of using AI to assist research into developing prevention measures for addictive behaviors.

Large Langage Models | Automatic Speech Recognition (ASR) | Texts | Audios

11:45-12:00
Lunch Poster

Pedro Macías Gordaliza, FBM,CHUV, Poster 19

Attribution Analysis of Performance Degradation in Medical Image Segmentation : *Our research investigates the attribution of performance degradation in medical image segmentation models when deployed on unseen datasets. Building upon established causal inference methodologies previously applied to tabular data, we extend these frameworks to address the complexities of high-dimensional medical imaging data. We propose a comprehensive approach to quantify how various factors in the annotation pipeline—including population shifts, imaging equipment variations, and annotator differences—contribute to performance drops. While traditional causal methods using interventions and counterfactuals provide valuable insights, we particularly focus on Importance Sampling (IS) techniques for their computational efficiency. This approach allows us to estimate distributional shifts in the high-dimensional space of medical images and their segmentation masks. Our methodology offers both theoretical rigor for the machine learning community and practical applicability for medical practitioners, providing interpretable attribution of performance degradation causes in clinical deployment scenarios.*

Random Forests | XGBoost | Convolutional neural networks | Diffusion models | Generative models | Tabulars | Images

Lunch Poster

Tanguy Falk, FBM,CHUV, Poster 1

Poster en lien avec mon travail de master (MLS - bioinformatics) portant sur l'extraction automatisée de données textuelles et visuelles en rapport à la mue des arthropodes en utilisant des outils tels que LLMs et CNNs. Les documents utilisés pour extraire les données sont principalement des articles scientifiques et des observations tirées de bases de données de sciences citoyennes telles que iNaturalist. J'effectue ce travail dans le groupe de Marc Robinson-Rechavi dans le département de la DEE.

Convolutional neural networks | Transformers | Large Langage Models | Tabulars | Images | Texts

Lunch Poster

Duy Cat Can, FBM,CHUV, Poster 3

Diagnosing Alzheimer's Disease (AD) using radiology images remains a complex challenge due to the need for both high predictive accuracy and clinical explainability. We present VisTA (Vision-Text Alignment Model), a multimodal framework that bridges this gap by aligning radiology images with expert-verified abnormalities and textual descriptions through contrastive learning. VisTA stands out by achieving state-of-the-art performance using only 170 fine-tuning samples, significantly outperforming large-scale pre-trained models like BiomedCLIP. The model offers a modular diagnostic pipeline that mirrors real-world clinical workflows : identifying abnormalities, retrieving similar reference cases, and providing interpretable explanations, all culminating in an evidence-driven AD diagnosis. VisTA achieves 74% accuracy (AUC : 0.87) in abnormality retrieval and 88% accuracy (AUC : 0.82) in dementia prediction, setting new benchmarks for explainable AI in healthcare. By integrating reliability, interpretability, and efficiency, VisTA advances machine learning's role in clinical decision-making.

Direct action neural networks | Transformers | Large Langage Models | Generative models | Explainable AI | Vision Transformer | Contrastive Learning | Multimodality | Zero-shot Prediction Images | Texts

Lunch Poster

Jacques Soutter, FGSE, Poster 4

Enhancing landslide prediction in Switzerland by combining the landslide historic and kilometer-scale hourly precipitation data over the past 20 years. Ensemble learning on classifiers.

Random forests | XGBoost | Direct action neural networks | Tabulars

Lunch Poster

Theo Maffei, FBM,CHUV, Poster 5

The heterogeneity of the tumor microenvironment reflects diverse and complex variations in the characteristics of the diverse cell populations that constitute a tumor, known to be tightly linked to disease evolution and treatment response. Advances in spatial omics techniques and computational methods enabled more reliable representation of the tumor microenvironment. However, the optimal approach for learning meaningful representations for patient stratification from a graph-based tumor model remains unclear. This study investigates how self-supervision in graph-based architectures might enhance the learning of meaningful representations for patient classification in a tumor modelling context. We evaluate the effectiveness of the self-supervised approach by benchmarking it against a baseline random forest classifier and a graph neural network architecture.

Linear/logistic regression | Graph neural networks | Images | Graphs

Lunch Poster

João Pedro, FBM,CHUV, Poster 7

The go/no-go reversal task can be modeled using reinforcement learning. An agent updates its policy via prediction error signals. When reward contingencies flip, rapid re-estimation of value functions and policy adjustments indicate efficient learning, measurable through error rate reduction and improved decision metrics. We test how different learning methods impact performance on the task.

Generative models | Tabulars | Reinforcement Learning

Lunch Poster

Léa Defferrard, FBM,CHUV, Poster 9

Spatial unilateral neglect (SUN) is a common post-stroke attentional deficit that severely impacts daily life. This study models SUN using reinforcement learning (RL) to analyze visuospatial behaviors and assess rehabilitation strategies. A virtual RL agent was trained to perform a visuospatial target detection task, generating behavioral metrics reflecting normal cognitive function. Simulated lesions were introduced by freezing neurons within the agent's neural network, replicating attention deficits observed in brain-injured patients. Machine learning techniques predicted lesion locations based on task performance, while dimensionality reduction helped identify affected neural functions. Rehabilitation strategies were tested within this framework to evaluate their effectiveness in restoring normal behavior. This approach enables precise deficit quantification and lesion-based therapy optimization, bridging computational neuroscience and clinical applications. By modeling SUN with RL, we provide insights into attentional control mechanisms and develop tailored rehabilitation strategies.

Reinforcement Learning | Auto data generation

Lunch Poster

Filippo Quarenghi, FGSE, Poster 10

This study compares different domain adaptation techniques to improve the generalization of downscaling models. Downscaling is crucial for enhancing the spatial and temporal resolution of weather and climate projections, facilitating more accurate regional impact assessments and supporting informed decision-making in agriculture, water management, and risk modelling. We focus on super-resolution techniques where low-resolution inputs are transformed into high-resolution outputs for variables like temperature, humidity, pressure, and precipitation. Specifically, we apply blurring to a 10-year, storm-resolving COSMO simulation over Europe, reducing the horizontal grid spacing to approximately 12 km before downscaling it back to the original 2.2 km resolution. To improve model generalization across unseen topographical conditions, we implement domain adaptation techniques evaluating the performance through tailored spatial evaluation metrics such as the Structure-Amplitude-Location (SAL) metric and the Fractions Skill Score (FSS). We hypothesize that domain adaptation will lead to superior generalization across diverse terrains, improved data efficiency, and better convergence, outperforming standard domain adaptation techniques such as Dropout and Batch Normalization. These models will ideally learn transferable mapping, allowing high-quality data from well-studied regions to enhance predictions in data-sparse areas. Additionally, we believe this approach may facilitate the discovery of more general relationships between mesoscale and storm-scale thermodynamic environments, improving the broader applicability of deep learning-based downscaling methods.

Convolutional neural networks | Diffusion models | Physics informed neural networks | Images | Climate data

Lunch Poster

Matteo Ferrarotti, FGSE, Poster 14

The cascading effect of wildfires on landslides represents a clear example of multi-hazard. Wildfires modify the local conditions of slopes in time, acting as a preparatory process for landslide triggering. A major challenge in studying post-wildfire landslides is the absence of comprehensive inventories documenting their occurrence. In this work, a numerical model was used for simulating various wildfire propagation scenarios, by setting different enucleation points and meteorological conditions. The results were used to generate quantitative predictions of expected shallow landslides over burned areas using a physically based model, PARSIFAL. The outputs of these analytical frameworks provide simulated post-wildfire landslides which have been used to train a Physics-Inspired Machine Learning model, along with several environmental predisposing factors for both wildfires and landslides. This tool can be useful for territorial planning and risk management, being capable of assessing post-wildfire landslide scenarios given by an ignition point, meteorological conditions, and predisposing factors.

Physics informed neural networks | Auto data generation

Lunch Poster

Peter Hilpert, SSP, Poster 20

In recent years, meta-researchers have provided numerous recommendations addressing questionable research practices concerning the testability of theories and the methods used to test them. These recommendations, while valuable, are often fragmented, making it challenging for applied researchers to understand how to integrate them into the complex research process. Additionally, there is the rarely discussed issue of coherence. Even if researchers correctly execute each step in the research process, this does not guarantee the necessary logical coherence between steps required to render the results interpretable. To investigate the prevalence of these problems in applied research, we systematically examined 100 publications from leading journals in psychology. In these manuscripts, we identified a multitude of problematic patterns, inconsistencies between research steps, and how errors at one point in the manuscript undermine subsequent stages. Overall, we highlight critical areas in the research process where quality and methodological rigor should be improved. Given that manual evaluation of all these problems in scientific manuscripts is highly resource-intensive, we are developing an automated system for detecting, extracting, and evaluating problems in research practices. The system relies on a large language model (LLM) that utilizes a Retrieval-Augmented Generation (RAG) system. This system will offer a scalable solution for addressing these problems in large research corpora. In sum, our findings not only shed light on the frequency and nature of such problems, including issues with coherence, but also demonstrate the potential of LLM and RAG-based tools to improve research transparency and theoretical alignment.

Lunch Poster

Pierre Heuzé, FBM,CHUV, Poster 21

Artificial intelligence and machine learning play a key role in scientific data analysis, predictive modeling and improved medical diagnosis. PACTT, the technology transfer office of CHUV and UNIL, promotes open-source publications where compatible, to accelerate knowledge sharing and innovation. It is also involved in projects linked to artificial intelligence and machine learning, given its role as a link between academic research and industry. PACTT facilitates the valorization of scientific discoveries and the transmission of academic knowledge, notably through licensing, collaboration with industry and the creation of start-ups. The table below shows the ideal way to pass on UNIL's knowledge, while preserving the various interests involved, both in terms of commercialization and community dissemination.

Lunch Poster

Haokun Liu, FGSE, Poster 22

Air pollution is a significant environmental hazard that poses serious risks to public health, making accurate modeling crucial for assessing human exposure. However, many studies rely on a limited number of monitoring stations and localized dispersion models, which may not capture the full complexity of air pollution distribution. To address these limitations, statistical and machine learning approaches are increasingly employed to model air pollution over broad spatial and temporal scales. In this study, we leveraged land use data, meteorological conditions, and topographical features to estimate the ground-level distribution of PM2.5 and NO2 across Switzerland. We developed and compared various statistical and machine learning models to generate high-resolution annual and monthly air pollution maps. These findings provide valuable insights for further research on the health risks and long-term effects associated with air pollution exposure, supporting more informed public health intervention strategies.

Linear/logistic regression | Random Forests | XGBoost | Chronological series | Geospatial data

14:00-14:15

Jaime Barranco Hernandez, FBM,CHUV, No poster

This study addresses the need for accurate 3D segmentation of the human eye and orbit from MRI to improve ophthalmic diagnostics. Past efforts focused on small sample sizes and varied imaging methods. Here, two techniques (atlas-based registration and supervised deep learning) are tested for automated segmentation on a large T1-weighted MRI dataset. Results show accurate segmentations of the lens, globe, optic nerve, rectus muscles, and fat. Additionally, the study automates the estimation of axial length, a key biomarker.

Convolutional nerual networks | Images

14:15-14:30

Andres Martinez Torres, SSP, No poster

The qualitative analysis of the discourse of communities is limited by factors such as data selection (necessary due to the size of the communities), researcher bias… Unsupervised topic modelling methods can analyse large datasets. However, they only provide word cooccurrence lists from which then 'topics' need to be abstracted and validated — mitigating but not solving the problems mentioned above and introducing other issues like dealing with the effects of data preprocessing. LLMs provide an opportunity to take a model that 'understands' language (through term associations), and see how it changes with 'exposure' (fine-tuning) with the language of a specific community. These models are not 'blank slates', as neither are humans entering communities — but they might be reproducible slates. This presentation explores the usage of LLMs to answer the question : How can Large-Language Models be used to better understand the idiosyncratic language of internet subcommunities ?

Large Langage Models | Texts

14:30-14:45
Lunch Poster

Kejdi Lleshi, FGSE, Poster 16

Poster Summary : *Cold episodes during the Quaternary have left visible glacial footprints on the landscapes of areas that were formerly glaciated. Climate proxies derived from these glacier footprints can be invaluable for reconstructing paleoclimate dynamics. However, most inversion methods are either 'manual' or heuristic and do not utilise the steepest gradient descent. With automatic differentiation, we have developed an invertible glacier evolution model. This model simulates the evolution of glaciers over a time series and stores all mathematical operations in a computation graph. It then uses chain derivation to compute the derivative of the cost function concerning the unknown climate parameters. The main limitation of this model is its memory usage. We employ AutoGrad's checkpoints to retain some crucial operations and subsequently recompute other operations during the backward pass. This new model allows us to converge more rapidly and to a superior solution compared to current models.* Presentation Summary : transformers/chatbots.

Automatic Differentiation (AD) | Chronological series

14:45-15:00

Roberto Miele, FGSE, No poster

Diffusion models (DMs) represent cutting-edge approach in computer vision, with successful scientific applications, such as atmospheric science and medical imaging. This work investigates DMs for parametrizing complex geological modeling and geophysical data inversion. In our approach, we represent a sedimentary scenario in a training dataset, consisting of channel bodies and corresponding acoustic impedances (IP), obtained from conceptual geological modeling and geostatistical simulations. The designed DM, trained on this dataset, demonstrates superior modeling accuracy compared to analogous Generative Adversarial Networks proposed for the same task, achieving up to 26% more accuracy in facies modeling and up to 90% more in reproducing IP distributions. Furthermore, we show that by implementing improved inversion techniques in diffusion modeling the trained network be used in Bayesian inversion of various geophysical data, and be conditioned on directly observed data (e.g., well logs or cores). This work represents an initial step in developing diffusion models for geophysics.

Diffusion models | Image generative models | Auto data generation

15:30-15:45

Maria Conticchio, FBM,CHUV, No poster

1.Current indications in the treatment of CRLM : the evalutation of more recent evidences about the accurate strategy of treatment between the standardization of care and the tailored therapy, 2.The application of a novel AI for the creation of a Decision Algorithm for Therapeutic Strategy in CLM, 3.Development of an Application for Algorithm Validation

Decision trees | Auto data generation

15:45-16:00

Kevin Borsos, FBM,CHUV, No poster

Cardiac magnetic resonance (CMR) imaging is a powerful tool for assessing both structural and functional cardiac health. Free-running CMR is a recently proposed technique that offers several advantages over conventional CMR imaging including being able to acquire data during free-breathing. However, the required compressed sensing (CS) reconstruction of this high-dimensional data is time-consuming, limiting its widespread clinical adoption. In this work, we propose a deep-learning based reconstruction to rapidly obtain free-running CMR images of high quality. A modified residual neural network is trained in a supervised manner on CS images to rapidly reconstruct free-running CMR data. Our deep learning approach provides comparable motion-resolved image quality to CS in less than one percent of the time required by CS. Our work demonstrates potential for rapid inline reconstruction of free-running CMR images for the first time, constituting a preliminary step towards bringing "single-click" free-breathing CMR to wider patient populations.

Convolutional neural networks | Images | Videos | Chronological series