VALPRED 2024

Prediction, validation and extremes

CAES du CNRS,24, rue du Coin
73500, Aussois, France
☎ (+33) 4 79 20 42 05

Topic

This workshop will deal with the general topic of forecasting with a particular emphazis on extremes. In a nutshell, our goal is to mix around 60 researchers, statisticians, meteorologists, energy providers and other end users who have a common interest in probabilisctic forecasting, especially for extreme events, and also in the assessment and validation of forecasts. The non-exhaustive list of potential topics will be: scoring rules theory, prediction methods and extreme value theory.

General information and important dates

The school will start on Monday 16th of December at lunch time (12:15) and will end on Thursday 19th of December (after lunch around 13h45). Registration, meals and lodging there will be covered during this period.
Check-in starts on Monday morning (front desk at the Aussois center).

It is possible to arrive on Sunday evening (not covered by the school) if you inform your time of arrival in the online form first. Please also provide the billing address needed for reimbursement in advance to reduce the administrative burden.
For local information regarding the center, contact
Mathilde Flechon (Mathilde.Flechon@caes.cnrs.fr)
Tel : +33 (0) 4 79 20 42 05 Fax : +33 (0) 4 79 20 30 44
or use the website *centre-paul-langevin*

For scientific information regarding the workshop, check the *VALPRED website*, or contact the local organising committee:
Clément Dombry (clement.dombry@univ-fcomte.fr),
Philippe Naveau (Philippe.Naveau@lsce.ipsl.fr),
Anne Sabourin (anne.sabourin@math.cnrs.fr),
Charles Tillier (charles.tillier@gmail.com) or
Olivier Wintenberger (olivier.wintenberger@sorbonne-universite.fr).

Transports

The Paul Langevin conference centre is located at an elevation of 1500 m, 7 km from the city of Modane with direct rail links to Paris and Turin. From the train station *Modane*, it is possible to take taxis to go from Modane to Aussois. For example, vans from the *Marius taxi* company can transport many participants coming from the same train (saving money and CO2 but you need to reserve them).

Meals

8:00-9:00 Breakfast
12:15-13:15 Lunch
19:15-20:00 Diner

Internet connections

Wifi code: centrepaullangevin73500

Talks

The talks will be 30 min long including questions. A projector and a whiteboard is at the disposal of the speakers. Prior to the session, speakers should upload their slides on the *collaborative folder here* or upload them on the conference computer via an usb key.

Institut Pierre Simon Laplace — Sciences de l'environnement
Fondation mathématique FMJH Jacques Hadamard
SORBONNE UNIVERSITÉ

Scientific schedule

The main room for the conference is "LA SCOLETTE". The poster session will take place in room "LA NORMA". Coffee breaks are served at the bar.

	MONDAY	TUESDAY	WEDNESDAY	THURSDAY
9 : 00		**Segers Johan** *Multivariate generalized Pareto distributions*	**Fildier Benjamin** *Different modes for the organization of atmospheric convection*	**Segers Johan** *Multivariate generalized Pareto distributions*
10 : 00		Coffee break	Coffee break	Coffee break
10 : 30		**Huet Nathan** *Multi-site modeling and reconstruction of past extreme sea levels* **Dell'Oro Lorenzo** *Flexible space-time models for environmental extreme data* **Cotsakis Ryan** *On the spatial extent of extreme threshold exceedances*	**Mayala Moria** *Predictive probability inference using random forests in the context of unbalanced classification* **Franchini Antoine** *Quantification of extrapolation performances of extreme quantile estimators* **Rizzeli Stefano** *Statistical prediction of Peaks Over Threshold*	**Baeriswyl Fabien** *Limiting results and large deviations of Poisson cluster processes* **Heranval Antoine** *Analyzing the dynamics of extreme events with marked point processes* **Kwasniok Franck** *Robust extreme value analysis by semiparametric modelling of the entire distribution range*
12 : 15	Lunch	Lunch	Lunch	Lunch
14 : 00	**Fildier Benjamin** *Different modes for the organization of atmospheric convection*	Informal time[1]	Informal time	
15 : 00	Break	Break	Break	
15 : 15	**Biegert Tobias** *Probabilistic Benchmarks and Post-Processing for Data-Driven Weather Forecasting* **Zimmermann Monika** *Spatial Meteorological, Socio-Economic, and Political Risks in Probabilistic Electricity Demand Forecasting*	Informal time	Informal time	
16 : 15	Coffee break	Coffee break	Coffee break	
16 : 30	**Dutta Ritabrata** *Minimum scoring rule for probabilistic forecasting* **Ólafsdóttir Helga Kristín** *Fast and robust scoring rule inference for spatial statistics* **Pic Romain** *Proper Scoring Rules for Multivariate Probabilistic Forecasts based on Aggregation and Transformation* **Poster presenters** *3 minute presentations of posters*	**Barbaux Occitane** *Designing life levels of Extreme Temperature by 2100* **Aka Samira** *Discrete Multivariate Generalized Pareto Distribution with application to dry spells* **Doize Antoine** *Stochastic Weather Generator for precipitation modelling, with focus on extreme rainfall and long dry spells* **Jacquemin Grégoire** *Estimating the return period of climate compound events using a non parametric bivariate Generalized Pareto representation*	**Mourahib Anas** *Extremal graphical models with non-standard extreme directions* **Lhaut Stephane** *Testing parametric models for the angular measure for bivariate extremes* **Durand Amaury** *Power comparison of sequential testing by betting procedures* **Hirsch Simon** *(Multivariate) Online Distributional Regression*	
19 : 15	Dinner	Dinner	Dinner	
20 : 00	*Poster session & welcome drink*	*Discussions*	*Fun time and karaoke*[2]	

1. Casual interactions among all participants are fundamental to built scientific and collaborative networking and to create new ideas. This is a time to chit-chat when hiking in the surroundings of Aussois or anything you like.
2. If you have any musical talents, feel free to bring your favorite instrument. Everyone will appreciate a little bit of music or singing.

Poster session

The poster session will take place on Monday evening 20:00 in room "LA NORMA" together with a welcome drink.

- **Adewoyin Rilwan**, *Joint Generalized Neural Models and Censored Spatial Copulas for Probabilistic Rainfall Forecasting*.
- **Drobac Nina**, *Signatures for Time Series Forecasting*.
- **Laveur Pearl**, *Can Gini measure of inequality discriminate between tail behaviors?*.
- **Mahomoreza Jean-Luc**, *Online Robust Time Series Forecasting Stochastic Covariates*.
- **Principato Guillaume**, *Conformal Predictions for Hierarchical Data*.
- **Tytgat Alexandre**, *Block-minima modeling of Antarctic's sea ice extent*.

BOOK OF ABSTRACTS

Abstracts for the mini-courses

Johan Segers, UC Louvain
Multivariate generalized Pareto distributions

The peaks-over-threshold approach for univariate extremes stipulates that the excess of a random variable over a high threshold can be modelled by the two-parameter family of generalized Pareto distributions. In statistical practice, the generalized Pareto distribution is fitted to the observed excesses of a variable over a high threshold, and the fitted model is used as a basis for extrapolation. For multivariate extremes, the same strategy can be followed using multivariate generalized Pareto distributions. The difference is now that these distributions no longer comprise a parametric family. In this mini-course, we will discover various surprising properties and representations of multivariate generalized Pareto distributions that will facilitate working with them for modelling multivariate extremes.

Benjamin Fildier, CNRS, École Normale Supérieure
Different modes for the organization of atmospheric convection

TBA.

Abstracts for Monday afternoon session

Tobias Biegert, Karlsruhe Institute of Technology
Probabilistic Benchmarks and Post-Processing for Data-Driven Weather Forecasting

In recent years, significant progress in machine learning technologies has enabled the development of various data-driven weather forecasting models, approaching, or even surpassing the skill level of physics-based numerical weather prediction (NWP) models. To facilitate comprehensive evaluation and comparison, WeatherBench 2, a recently developed benchmarking framework for data-driven weather prediction, incorporates predictions from a wide range of state-of-the-art AI models as well as traditional NWP models. However, despite these advancements, several important questions remain open. Most data-driven models primarily focus on deterministic point forecasts and lack the capability to generate probabilistic predictions, which, however, is crucial for optimal decision making and quantifying weather risk in applications. Further, while it has been widely demonstrated that physics-based NWP models substantially benefit from post-processing methods, which aim to correct systematic errors, the use of post-processing for data-driven weather models has not been explored in detail. Our overarching aim thus is to investigate the application of various post-processing techniques to potentially improve predictions, as well as to generate probabilistic forecasts from state-of-the-art data-driven as well as NWP model outputs. On the one hand, we thereby assess whether AI-based weather models benefit from post-processing in a similar manner to physics-based NWP models, and, on the other hand, enable a fair comparison of post-processed AI-based and post-processed physics-based models. We apply a range of established statistical and machine learning-based post-processing methods to predictions provided within the WeatherBench 2 framework, for the eight variables defined as headline scores (Z500, T850, Q700, WV850, T2M, WS10, MSLP, TP24hr), and systematically evaluate the effectiveness of these methods for improving deterministic and probabilistic forecasts. Our evaluation reveals that post-processed probabilistic forecasts can outperform the ensemble predictions from the European Centre for Medium-Range Weather Forecasts for shorter lead times of up to one week for selected variables, but the results tend to vary widely across variables, lead times and forecasting models.

Monika Zimmerman, University Duisburg Essen
Spatial Meteorological, Socio-Economic, and Political Risks in Probabilistic Electricity Demand Forecasting

Unlike traditional commodities, electricity is difficult to store and, thus, delivered via a cross-country interconnected grid that requires a continuous balance between supply and demand. This characteristic necessitates high-resolution mid-term electricity demand forecasts to inform financial and operational decisions. The primary challenge in creating these forecasts lies in disentangling the uncertain, spatially varying, and cross-country-dependent impacts on load. Additionally, the non-stationarity of socio-economic and political load effects adds complexity. To address these challenges, we present a novel forecasting method for mid-term hourly electricity demand that is probabilistic and multivariate across 24 European countries. This approach combines multivariate simulated socio-economic, political and temperature trajectories that account for climate change within Generalized Additive Models (GAMs). Built as interpretable additive model in smooth effects, it provides practitioners with a tool to track the hourly load impact of each input at gigawatt precision. We evaluate our method using over 9 years of data (2015-2024), demonstrating significantly improved CRPS compared to standard benchmarks. Our results indicate that a single Pan-European socio-economic and political trend explains much of the non-stationarity in load across Europe. Risk scenario analyses highlight the vulnerability of countries dependent on electric heating during extreme weather events, underscoring the need for probabilistic, cross-country dependent forecasting as electric heating becomes more prevalent across Europe.

Ritabrata Dutta, University of Warwick
Minimum scoring rule for probabilistic forecasting
Scoring rules (eg. Energy score, CRPS) have been used long to assess the accuracy of forecasts. In our works we have used scoring rules to train the models instead, providing a better predictive performance. Our developed framework envelops traditional classical and Bayesian statistics, but this setup is also amenable to likelihood-free scenarios and scalable to high dimensional parameter space. To train weather forecasting models, we further proposed prequential scoring rules and show their consistent performance using both theory and simulations. In the talk, I will also highlight how this framework can be extended for forecasting weather for many time points in future jointly using a signature kernel score or adapted for models of extreme values.
(Joint work with Lorenzo Pacchiardi, Sherman Khoo, Peter Dueben, Rilwan Adewoyin, Shreya Sinha Roy, Francesca Bassini, Yuehao Xu, Andreas Futschik, Philippe Naveau.)

Helga Kristín Ólafsdóttir, Chalmers University of Technology
Fast and robust scoring rule inference for spatial statistics
Scoring rules are not only useful for model evaluation, but have also successfully been used for model inference. We propose parameter inference of spatial models by optimising the predictive ability through the leave-one-out cross-validation score (LOOS), instead of maximising the likelihood. This approach is studied for different Gaussian spatial models. For models with sparse covariance matrises, the LOOS approach results in fast computations compared to the likelihood approach. Moreover, the LOOS approach allows affecting the robustness to outliers and sensitivity to stationarity. Applying the estimation methods on ERA5 temperature reanalysis data for the average July temperature in contiguous United States for years 1940 to 2023, showed that the LOOS estimator improved predictive performance in a fraction of the computation time.

Romain Pic, Université de Genève
Proper Scoring Rules for Multivariate Probabilistic Forecasts based on Aggregation and Transformation
Proper scoring rules are an essential tool to assess the predictive performance of probabilistic forecasts. However, propriety alone does not ensure an informative characterization of predictive performance and it is recommended to compare forecasts using multiple scoring rules. With that in mind, interpretable scoring rules providing complementary information are necessary. We formalize a framework based on aggregation and transformation to build interpretable multivariate proper scoring rules. Aggregation-and-transformation-based scoring rules can target application-specific features of probabilistic forecasts, which improves the characterization of the predictive performance. This framework is illustrated through examples taken from the weather forecasting literature and numerical experiments are used to showcase its benefits in a controlled setting. Additionally, the framework is tested on real-world data of postprocessed wind speed forecasts over central Europe. In particular, it is shown that it can help to bridge the gap between proper scoring rules and spatial verification tools.

Abstracts for Monday evening poster session

Rilwan Adewoyin, ECMWF
Joint Generalized Neural Models and Censored Spatial Copulas for Probabilistic Rainfall Forecasting
This work introduces a novel approach for generating conditional probabilistic rainfall forecasts with temporal and spatial dependence. A two-step procedure is employed. Firstly, marginal location-specific distributions are jointly modelled. Secondly, a spatial dependency structure is learned to ensure spatial coherence among these distributions. To learn marginal distributions over rainfall values, we introduce joint generalised neural models which expand generalised linear models with a deep neural network to parameterise distributions. To understand the spatial dependency structure, a censored latent Gaussian copula model is presented and trained via scoring rules. Leveraging the underlying spatial structure, we construct a distance matrix between locations, transformed into a covariance matrix by a Gaussian Process Kernel depending on a small set of parameters. To estimate these parameters, we propose a general framework for the estimation of Gaussian copulas employing scoring rules as a measure of divergence between distributions. Uniting our two contributions, namely the joint generalised neural model and the censored latent Gaussian copulas into a single model, our probabilistic approach generates forecasts on short to long-term durations, suitable for locations outside the training set. We demonstrate its efficacy using a large UK rainfall data set, outperforming existing methods

Nina Drobac, Sorbonne Université
Signatures for Time Series Forecasting
TBA.

Pearl Laveur, Université Grenoble Alpes
Can Gini measure of inequality discriminate between tail behaviors?
The richest 1% on the planet own almost half of the world's wealth, according to a study by Oxfam International published in 2021. This marked inequality in the distribution of wealth is of major concern to our society and economy. Several measures are known to quantify inequality in the distribution of a variable of interest, the Gini index being the most popular. However, the literature suggests that this inequality measure is unable to discriminate between different tail behaviours. For instance, heavy-tailed and short-tailed distributions may have the same Gini index. In this work, we rigorously state that the Gini index does not discriminate between maximum domains of attraction. To overcome this drawback, an "extreme" Gini index is introduced, based on the excesses over a high threshold. We show that, when the threshold tends towards the endpoint, this extreme Gini index is able to discriminate between different distribution tails.

Jean-Luc Mahomoreza, Sorbonne Université
Online Robust Time Series Forecasting Stochastic Covariates
The Kalman-Bucy filter is used to estimate and predict state-space models. It frequently faces uncertain observations, such as those derived from noisy measurements of physical signals (including sensors, weather data, etc.). Instead of considering covariates as deterministic quantities, as is conventionally done, our purpose is to take their uncertainties into account and address a robust filter. More specifically, the Kalman filter is a Bayesian approach that estimates the posterior distribution of the state conditionally on past observations. We refine this approach by integrating the uncertainties of the covariates into the model, i.e., considering them as stochastic variables whose distributions need to be estimated.

Guillaume Principato, Université Paris Saclay
Conformal Predictions for Hierarchical Data
Reconciliation has become an essential tool in multivariate point forecasting for hierarchical time series. However, there is still a lack of understanding of the theoretical properties of probabilistic Forecast Reconciliation techniques. Meanwhile, Conformal Prediction is a general framework with growing appeal that provides prediction sets with probabilistic guarantees in finite sample. In this paper, we propose a first step towards combining Conformal Prediction and Forecast Reconciliation by analyzing how including a reconciliation step in the Split Conformal Prediction (SCP) procedure enhances the resulting prediction sets. In particular, we show that the validity granted by SCP remains while improving the efficiency of the prediction sets. We also advocate a variation of the theoretical procedure for practical use. Finally, we illustrate these results with simulations.

Alexandre Tytgat, UCLouvain
Block-minima modeling of Antarctic's sea ice extent
The Antarctic sea ice extent (SIE) has exhibited unusual variability in recent years, including record lows in annual minima observed in 2017, 2022, 2023, and 2024. This marks a sharp contrast with the sea ice extent evolution between 1979 and 2015, during which yearly mean sea ice extent slightly increased despite global warming. Such events raise critical questions about their probabilities under current and future conditions, as well as the potential factors driving such events. To address these questions, we use block-minima models to analyze and interpret the yearly minima of the Antarctic SIE. Additionally, we extend this approach to the five regional basins that constitute the Antarctic, aiming to shed some light on their contributions to the recent extreme events. Our results suggest that the trend in SIE yearly minima is best captured by a piecewise linear model for the mean, and an increasing scale parameter over time, and that the recent record lows, though unprecedented, were not highly improbable given recent trends. The regional analysis indicates that the Ross Sea, Weddell Sea and Western Pacific Ocean basins are the primary contributors to the recent decline in total SIE.

Abstracts for Tuesday morning session

Nathan Huet, Telecom Paris
Multi-site modeling and reconstruction of past extreme sea levels along the Atlantic French coast
Appropriate modeling of extreme sea levels is crucial, particularly for coastal risk management. Our study focuses on modeling extreme sea levels along the French Atlantic coast, with a particular emphasis on investigating the extremal dependence structure between stations. We employ the peak-over-threshold framework, where a multivariate extreme event is defined whenever at least one location records a large value, though not necessarily all stations simultaneously. A novel method for determining an appropriate level (threshold) above which observations can be classified as extreme is proposed. We investigate two joint approaches separately. First, we utilize the Multivariate Generalized Pareto Distribution (MGPD) to model these extremes, which is well-suited for this type of analysis. Given the beneficial properties of this distribution, we derive a generative model for extreme sea levels at a station conditionally based on extreme values at other nearby stations. Second, we examine the point estimate performance of a novel extreme regression model. This specific regression framework enables accurate point predictions using only the 'angle' of input variables, i.e., input variables divided by their norms. The ultimate objective of this work is to reduce the uncertainty behind risk

management quantities (e.g., return values) at stations with limited historical records by incorporating new data based on extreme sea levels from stations with longer records, such as Brest and Saint-Nazaire, which have over 150 years of records.
Joint work with Philippe Naveau and Anne Sabourin.

Lorenzo Dell'Oro, Università di Padova
Flexible space-time models for environmental extreme data
Extreme value analysis is critical for understanding rare and extreme events, whose study is of significant interest in various fields, notably in environmental sciences. The most common modeling approach for spatio-temporal extremes relies on asymptotic models, such as max-stable and r-Pareto processes, respectively for block maxima and peaks over high thresholds. However, the lack of flexibility that characterizes the dependence structure of these models makes them unable to capture the empirical pattern of extreme events becoming more localized in space and time as their severity increases. To address this limitation, some authors have recently focused on models capable of flexibly describing the "sub-asymptotic" dependence of data, many of which utilize a random scale construction. This study presents an extension of the spatial random scale mixture model proposed by Huser and Wadsworth (2019, JASA) to the spatio-temporal domain, providing a comprehensive framework for characterizing the dependence structure of extreme events across both dimensions. Indeed, the model is able, through parametric inference, to discriminate between asymptotic dependence and independence, concurrently both in space and time. Due to the high complexity of the likelihood function for the proposed model, parameter estimation relies on a simulation approach based on neural networks, which leverages summaries of the sub-asymptotic dependence present in the data. The effectiveness of the model in assessing the limiting dependence structure of spatio-temporal processes is shown both via simulation studies and through an application to rainfall datasets.

Ryan Cotsakis, Université de Lausanne
On the spatial extent of extreme threshold exceedances
We introduce the extremal range, a local statistic for studying the spatial extent of extreme events in random fields on Rd. Conditioned on exceedance of a high threshold at a location s, the extremal range at s is the random variable defined as the smallest distance from $s \in R^d$ to a location where there is a nonexceedance. We leverage tools from excursion-set theory, such as Lipschitz- Killing curvatures, to express distributional properties of the extremal range, including asymptotics for small distances and high thresholds. The extremal range captures the rate at which the spatial extent of conditional extreme events scales for increasingly high thresholds, and we relate its distributional properties with the well-known bivariate tail dependence coefficient and the extremal index of time series in Extreme-Value Theory. We calculate theoretical extremal-range properties for commonly used models, such as Gaussian or regularly varying random fields. Numerical studies illustrate that, when the extremal range is estimated from discretized excursion sets observed on compact observation windows, the distribution of the resulting estimators appropriately reproduces the theoretically derived links with the Lipschitz- Killing curvature densities.
arxiv preprint https://arxiv.org/abs/2411.02399

Abstracts for Tuesday afternoon session

Barbaux Occitane, CNRM
Designing life levels of Extreme Temperature by 2100
"In the context of climate change, design levels for environmental extremes have to be redefined. The objectives of this work are twofold. The first one is to provide a single indicator that summarizes relevant and interpretable information about large values in times series, even when they cannot be assumed stationary. A second objective is to capture all inferential uncertainties for this single indicator.
The usual risk indicator, the annual return level, is only defined in a stationary context. Since it is defined as a level corresponding to an annual probability of excess, a singular value cannot be defined for a period of interest. It is necessary to select another risk indicator able to cover an entire time period. The Equivalent Reliability, which defines the probability of the maximum event during the period, was chosen because its probability of excess over a period is the same as the total probability of excess over the same period for a stationary return level (Liang 2016 et Hu 2018). For our second objective, we adapt the predictive distribution to Equivalent Reliability to include uncertainty over the parameters' full distribution. This method accounts for a wide range of possible extreme values while providing a unique value fit for design calculation.
Our approach was applied to annual temperature maxima at a single location in the South of France for various emission scenarios. Estimation is done using the Bayesian statistical framework proposed by Robin and Ribes (2020), with a non-stationary generalized extreme value distribution that combines 26 CMIP6 general circulation models (GCMs) with an observational record.

Samira Aka, Square Management & LSCE
Discrete Multivariate Generalized Pareto Distribution with application to dry spells
Through theoretical results and illustrative examples, this presentation settles the construction and properties of MDGPD. Additionally, it provides practical insights into simulation techniques and data fitting approaches for MDGPDs. The proposed framework expands the toolkit for modeling extreme events in various domains, offering a methodology for analyzing multivariate discrete data with extreme values. An application of droughts will be presented.

Antoine Doizé, LSCE

Stochastic Weather Generator for precipitation modelling, with focus on extreme rainfall and long dry spells
TBA.

Grégoire Jacquemin, Mine Paris
Estimating the return period of climate compound events using a non parametric bivariate Generalized Pareto representation
Compound events, commonly defined as the "combination of multiple drivers and/or hazards that contributes to societal or environmental risk", often result in amplified impacts compared to individual hazards. This study examines compound flooding triggered by accumulated precipitation across two watersheds. In the general framework of multivariate extremal value theory, two models are explored and compared: (1) a copula-based approach that models joint exceedances over thresholds, and (2) a novel non-parametric approach employing Multivariate Generalized Pareto (MGP) distributions. Special attention is given to the treatment of strongly correlated data and non-concomitant compound events. The return period for such bivariate events is carefully defined and closed formula are obtained for both methodologies. Simulations reveal the MGP approach as particularly effective for strongly correlated variables. The methods are subsequently applied to ERA5 reanalysis data to analyze compound flooding scenarios, highlighting their practical implications for risk assessment and management.

Abstracts for Wednesday morning session

Moria Mayala, Sorbonne Université
Predictive probability inference using random forests in the context of unbalanced classification
TBA.

Antoine Franchini
Quantification of extrapolation performances on extreme quantile estimators
TBA.

Stefano Rizzeli, Universitá di Padova
Statistical prediction of Peaks Over Threshold: predictive density estimation, risk assessment and uncertainty quantification
In many applied fields it is desired to make predictions with the aim of assessing the plausibility of more severe events than those already recorded, to safeguard against calamities that have not yet occurred. This problem can be analysed using extreme value theory. We consider the popular peaks over a threshold method and show that the generalised Pareto approximation of the true predictive density of a future peak can be very accurate. We discuss both frequentist and Bayesian approaches for the estimation of such predictive density. We show the asymptotic accuracy of the corresponding estimators and, more importantly, prove that the resulting predictive inference is asymptotically reliable. We conclude discussing extensions of the proposed approach, e.g. to account for hetheroscedasticity and covariate information, and drawing a parallel with the problem of predicting future periodic maxima.

Abstracts for Wednesday afternoon session

Anas Mourahib, UCLouvain
Extremal graphical models with non-standard extreme directions
Graphical models for multivariate generalized Pareto distributions have received significant attention, especially when the exponent measure does not allocate mass on lower-dimensional extreme directions. Various models have been proposed in this context, exhibiting differing degrees of graphical complexity. For instance, while the logistic multivariate generalized Pareto distribution presents a trivial graphical structure, others entail more intricate representations, as demonstrated in recent literature. By utilizing certain models proposed in the literature to construct multivariate generalized Pareto distributions with lower-dimensional extreme directions, and building upon established theoretical frameworks of graphical conditional independence, we introduce a novel approach to construct non-standard extremal graphical models. These models encapsulate the underlying graphs of multivariate generalized Pareto distributions with lower-dimensional extreme directions. Subsequently, we employ this approach to extend existing graphical models, such as the extended logistic graphical model and the extended Hüsler—Reiss model.

Stéphane Lhaut, UCLouvain
Testing parametric models for the angular measure for bivariate extremes
The angular measure on the unit sphere characterizes the first-order dependence structure of the components of a random vector in extreme regions and is defined in terms of standardized margins. Its statistical recovery is an important step in learning problems involving observations far away from the center. In this paper, we test the goodness-of-fit of a given parametric model to the extremal dependence structure of a bivariate random sample. The proposed test statistic consists of a weighted L1-Wasserstein distance between a nonparametric, rank-based estimator of the true angular measure obtained by maximizing a Euclidean likelihood on the one hand, and a parametric estimator of the angular measure on the other hand. The asymptotic distribution of the test statistic under the null hypothesis is derived and is used to obtain critical values for the proposed testing procedure via a parametric bootstrap. Consistency of the bootstrap algorithm is proved. A simulation study illustrates the finite-sample performance of the test for the logistic and Hüsler–Reiss models. We apply the method to test for the Hüsler–Reiss model in the context of river discharge data.

Amaury Durand, EDF & Telecom Paris
Power comparison of sequential testing by betting procedures
TBA.

Simon Hirsh, Statkraft & University Duisburg-Essen
(Multivariate) Online Distributional Regression
Large scale streaming data are common in many modern ML applications and especially in the energy sector, where renewable production forecasts, meteorological data and price are updated continuously. At the same time, the rise of probabilistic forecasting yields the need for effective, incremental learning of the conditional distribution of the data. Against this background, we present regularized online estimation methods for (multivariate) distributional regression. Our method builds on the penalized iteratively reweighted least squares approach for the estimation of GAMLSS models and leverages online coordinate descent methods for parameter estimation. The talk will discuss univariate and multivariate approaches. For univariate distributional modelling, we provide simulation results and a forecasting study on electricity markets, showing the competitive performance in terms of forecasting accuracy, while reducing the computational efforts by 2-3 orders of magnitude. While the estimation principle translates seamlessly to multivariate distributions, a major complication is the quadratically growing model size. In current work in progress, we propose a possible remedy by either using AD-r regularized, Cholesky-decomposition based parameterizations or by using low-rank approximations of the covariance matrix, both in combination with path-based estimation procedures with early stopping to keep model complexity manageable. Initial, small-scale simulation results seem promising in terms of accuracy and computational effort.

Abstracts for Thursday morning session

Fabien Baeriswyl, Université de Lausanne
Limiting results and large deviations of Poisson cluster processes
In this talk, we focus on the hidden regular variations of certain Poisson cluster marked point processes. The hidden regular variations of these processes are expressed as successive convergences on a space of point measures, where increasingly broad cones of point measures are successively removed. We highlight that these principles of successive hidden regular variations are equivalent to classical large deviation principles, and the rate function is derived using these tools. From these results concerning point processes, we deduce results (on the Skorokhod space) of hidden regular variations for certain functionals.
This work is a collaboration with Olivier Wintenberger.

Antoine Heranval, INRAE
Analyzing the dynamics of extreme events with marked point processes
The aim is to merge probabilistic methods for marked point processes, extreme events, and spatial graphs in order to propose a new framework for better understanding spatial and temporal interactions between different types of extreme events in the climate system. Each extreme event (whether temperature, precipitation, wind speed, atmospheric pressure, etc.) for a specific climate variable and region will be represented by a dot indicating the time of its occurrence, as well as by one or more markers describing key characteristics of the episode, such as maximum, spatial extent, duration, accumulation, etc. First- and second-order statistical characteristics and second-order statistical features of point processes, such as the intensity function and various pair and mark correlation functions, will be adapted to study the multiscale structures of the set of marked points. New classes of models are also developed for marked point processes, defined on a spatial graph representing the geographical proximity between regions, in order to simulate new realizations of the set of such events.

Franck Kwasniok, University of Exeter
Robust extreme value analysis by semiparametric modelling of the entire distribution range
Traditional extreme value analysis based on the generalised extreme value (GEV) or the generalised Pareto distribution (GPD) suffers from two drawbacks: (i) Both methods are wasteful of data as only block maxima or exceedances over a high threshold are used and the bulk of the data is disregarded, resulting in a large uncertainty in the tail inference. (ii) In the peak-over-threshold approach the choice of the threshold is often difficult in practice as there are no really objective underlying criteria. Here, two approaches based on maximum likelihood estimation are introduced which simultaneously model the whole distribution range and thus constrain the tail inference by information from the bulk data. Firstly, the bulk matching method models the bulk of the distribution with a flexible exponential family model using polynomial, spline or trigonometric basis functions, and the tail with a GPD. The two distributions are linked together at the threshold with appropriate matching conditions. The method becomes insensitive to the choice of the threshold. Secondly, in the extended generalised Pareto distribution (EGPD) model for non-negative variables the whole distribution is modelled with a GPD overlaid with a transition probability density which is again represented by an exponential family. Appropriate conditions ensure that the model is in accordance with extreme value theory both for the lower and upper tail of the distribution. This approach completely avoids the choice of a threshold. The methods are successfully exemplified on simulated data as well as wind speed and precipitation data. Compared to the classical GEV and GPD analyses, the new techniques provide more accurate estimates of extreme quantiles with a smaller uncertainty associated to them.