VALPRED 2023

Workshop on (probabilistic) forecast and its verification
Validation et prédictions probabilistes

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

Topic

This workshop will deal with the topic of validation of ensemble forecasting. In a nutshell, our goal is to mix around 40 researchers, statisticians, meteorologists, energy providers and other end users who have a common interest in probability forecasting, especially in the assessment of different forecasts. The non-exhaustive list of potential topics will be: scoring rules theory, forecasting methods, extreme value theory, decision theory.

General information and important dates

The school will start on Monday 3rd of April at lunch time (12:30) and will end on Thursday 6th of april after lunch (13:30). Registration, meals and lodging there will be covered during this period. Early arrival on Sunday 2nd of April is available upon request. *Check-in at front desk at the Aussois CAES center starts on Sunday evening or on Monday morning*.

For local information regarding the center, contact Christa Balzer (christa.balzer@caes.cnrs.fr, tel : +33 (0) 4 79 20 42 05) or use the website *centre-paul-langevin*

For scientific information regarding the workshop, check the *VALPRED website* , or contact the local organising committee: Charles Tillier (charles.tillier@gmail.com), Clément Dombry (clement.dombry@univ-fcomte.fr), Philippe Naveau (Philippe.Naveau@lsce.ipsl.fr) or Olivier Wintenberger (olivier.wintenberger@sorbonne-.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
12:30-13:30 Lunch
19:15-20:00 Dinner

Internet connections

Wifi code: centrepaullangevin73500

Sciences de l'environnement Institut Pierre Simon Laplace — Fondation mathématique FMJH Jacques Hadamard — SORBONNE UNIVERSITÉ — UBFC UNIVERSITÉ BOURGOGNE FRANCHE-COMTÉ

Scientific program

The workshop will take place in the conference room "LA SCOLETTE".

	MONDAY	TUESDAY	WEDNESDAY	THURSDAY
9 : 00		*Mini-course* **Benedikt Schulz** *Statistical and Machine Learning Methods for Postprocessing Ensemble Weather Forecasts*	*Mini-course* **Benedikt Schulz** *Statistical and Machine Learning Methods for Postprocessing Ensemble Weather Forecasts*	
10 : 00		Coffee break	Coffee break	10:00-12:00
10 : 30		*Time Series* **Amandine Pierrot** *On tracking varying bounds when forecasting bounded time series* **Björn Sonnenschein** *Probabilistic Intraday Wastewater Treatment Plant Inflow Forecast Utilizing Rain Forecast Data and External Flow Measurements* **Gloria Buritica** *On blocks estimators for cluster inference of heavy-tailed time series*	*Machine Learning* **Nicolas Lafon** *A variational auto-encoder approach to sample multivariate extremes* **Moria Mayala** *Subsampled base learners in extreme regions* **Thibault Modeste** *RKHS and some applications*	*Forecast Evaluation* **Sam Allen** *Asymmetric proper scoring rules* **Harris Nkuiate** *Evaluating Probabilistic Forecasts in the Presence of Observation Error* **Amaury Durand** *Sequential evaluation of a forecaster by betting* **Romain Pic** *Distributional forecasting and its evaluation with the CRPS*
12 : 30	Lunch	Lunch	Lunch	Lunch
15 : 30	*Extremes* **Marco Oesting** *Extremes in High Dimensions: Methods and Scalable Algorithms*	Informal time[1]	Informal time[1]	
16 : 30	Coffee Break	Coffee Break	Coffee Break	
17 : 00	*Extremes* **Fabien Baeriswyl** *Precise large deviations for certain cluster processes* **Nathan Huet** *Extremes of functional data: low dimensional representation via Karhunen-Loève expansion* **Paula Gonzalez** *A statistical method to model non-stationarity in precipitation records changes*	*Forecast applications* **Nicolas Eckert** *An overview of a few statistical issues in snow avalanche short and long-term forecasting* **Yannig Goude** *Adaptive probabilistic forecasting of net electricity demand* **Thomas Muschinski** *Robust weather-adaptive postprocessing using MOS random forests*	*Environmental applications* **Jonathan Koh** *Predicting risks of temperature extremes using large-scale circulation patterns with r–Pareto processes* **Juliette Legrand** *Bayesian spatio-temporal modelling of wildfire occurrences and sizes for climate model projections* **Léo Pfitzner** *Temperature forecasting with expert aggregation*	
19 : 15	Dinner	Dinner	Dinner	
20 : 00	Poster session + cocktail	*Fun time*[2]	*Karaoké*[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.

Detailed program

Tuesday, Wednesday 09:00-10:00 – Mini-course

Sebastian Lerch and Benedikt Schulz: Statistical and Machine Learning Methods for Postprocessing Ensemble Weather Forecasts

Nowadays, weather prediction is based on numerical models of the physics of the atmosphere. These models are usually run multiple times based on randomly perturbed initial conditions. The resulting so-called ensemble forecasts represent distinct scenarios of the future and provide probabilistic projections. However, these forecasts are subject to systematic errors such as biases and they are often unable to quantify the forecast uncertainty adequately. Statistical postprocessing methods aim to exploit structure in past pairs of forecasts and observations to correct these errors when applied to future forecasts. Over the past two decades, a variety of methods has been developed and postprocessing has become a standard practice in research and operations.

Monday 15:30-16:30 – Extremes

Marco Oesting: Extremes in High Dimensions – Methods and Scalable Algorithms

Extreme value theory has been explored in considerable detail for univariate and low-dimensional observations, but the field is still in an early stage regarding high-dimensional multivariate observations. In this paper, we focus on Hüsler-Reiss models and their domain of attraction, a popular class of models for multivariate extremes that exhibit some similarities to multivariate Gaussian distributions. We devise novel estimators for the parameters of this model based on score matching and equip these estimators with state-of-the-art theories for high-dimensional settings and with exceptionally scalable algorithms. We perform a simulation study to demonstrate that the estimators can estimate a large number of parameters reliably and fast; for example, we show that Hüsler-Reiss models with thousands of parameters can be fitted within a couple of minutes on a standard laptop. Their usefulness for applications is illustrated in a real data example on weather extremes.

Monday 17:00-18:30 – Extremes (chair C. Dombry)

Fabien Baeriswyl: Precise large deviations for certain cluster processes

In this talk, we propose certain tail asymptotics of cluster functionals (principally the maximum and the sum) for mainly two Poisson cluster processes, namely the renewal Poisson cluster process and the Hawkes process. We use these asymptotic results to derive precise large deviations principles for the mentioned cluster processes. Joint work with V. Chavez-Demoulin and O. Wintenberger.

Nathan Huet: Extremes of functional data: low dimensional representation via Karhunen-Loève expansion

The increasing availability of data of functional nature associated with the deployment of sensors providing massive measurements sampled at an ever finer granularity opens new perspective for unsupervised learning regarding extremes of functional data. We consider regularly varying stochastic processes with paths in $L^2([0,1])$, the Hilbert space of square integrable functions on $[0,1]$. The concept of Karhunen-Loève expansion, at the very heart of the $L^2([0,1])$-stochastic processes theory, is then investigated with the objective of recovering the development of a limit extreme process based on a sample of i.i.d. functional observations of a heavy-tailed process.

Paula Gonzalez: A statistical method to model non-stationarity in precipitation records changes

In this work we propose an extreme event attribution (EEA) methodology to analyze yearly maxima records, key indicators of climate change that spark media attention and research in the EEA community. Although they deserve a specific statistical treatment, algorithms tailored to record analysis are lacking. This is particularly true in a non-stationarity context. This work aims at filling this methodological gap by focusing on records in transient climate simulations. We apply our methodology to study records of yearly maxima of daily precipitation issued from the numerical climate model IPSL-CM6A-LR. Illustrating our approach with decadal records, we detect in 2023 a clear human induced signal in half of the globe, with probability mostly increasing, but decreasing in the south and north Atlantic oceans.

Tuesday 10:30-12:00 – Time series (chair P. Naveau)

Amandine Pierrot: On tracking varying bounds when forecasting bounded time series

We consider a new framework where a continuous, though bounded, random variable has unobserved bounds that vary over time. In the context of univariate time series, we look at the bounds as parameters of the distribution of the bounded random variable. We introduce an extended log-likelihood estimation and design algorithms to track the bound through online maximum likelihood estimation. Since the resulting optimization problem is not convex, we make use of recent theoretical results on Normalized Gradient Descent (NGD) for quasiconvex optimization, to eventually derive an Online Normalized Gradient Descent algorithm. We illustrate and discuss the workings of our approach based on both simulation studies and a real-world wind power forecasting problem.

Björn Sonnenschein: Probabilistic Intraday Wastewater Treatment Plant Inflow Forecast Utilizing Rain Forecast Data and External Flow Measurements

We present a seasonal probabilistic time series model for modelling short-term wastewater treatment plant inflow. To ensure suitability for practical implementation, the unconstrained parameters of the predictive distribution are modelled as linear functions of the input variables and non-linear effects caused by the wastewater network hydrodynamics are approximated by linear splines. In addition to water level measurements from within the sewer network and rain rate measurements, rain forecasts are incorporated as exogenous regressors, where historical rain forecasts are used for model calibration. The model performance is evaluated on historical data from a German wastewater treatment plant using deterministic and probabilistic scoring rules. Our results show that the model's prediction accuracy is comparable to an LSTM while providing better explainability to practitioners, but accurate real-time rain forecasts are mandatory for successful real-world implementation.

Gloria Buritica: On blocks estimators for cluster inference of heavy-tailed time series

In the framework of heavy-tailed time series, extremal observations cluster: an extreme value triggers a short period with numerous large observations. This behaviour is known to perturb classical inference procedures tailored for independent observations like high quantile inference. We aim to infer properties of the clustering effect by applying functions to consecutive observations with extremal behaviour. We recover classical statistics like the extremal index and cluster size probabilities with cluster inference. In this talk, we discuss the asymptotics of block estimators for cluster inference based on consecutive observations with large $l\alpha - norm$, where $\alpha > 0$ is the tail index of the series. Interestingly, in the case of ARMA models, our computations show that many cluster statistics have null asymptotic variance, as first conjectured in Hsing T. (1996).

Tuesday 17:00-18:30 – Forecast applications (chair O. Wintenberger)

Nicolas Eckert: An overview of a few statistical issues in snow avalanche short and long-term forecasting

In many mountain environments of the world, snow avalanches are a recurrent danger in winter, resulting in direct (casualties, damages to buildings, infrastructures, forests) and indirect (e.g. road closures that isolates high valleys from the rest of the world) losses. Risk management involves short and long term forecasting of avalanche hazard and its combination with elements at risk and their vulnerability. A specific difficulty is that mountain socio-ecosystems are now changing faster than ever, making snow avalanche risk highly non-stationary in all its components. This talk will provide a brief overview of some remaining statistical issues still encountered in the short and long term forecasting of snow avalanche risks. As statisticians working in this specific field are currently not numerous, the idea is to potentially involve more people, so as to, in the future, foster the development of innovative risk assessment and mitigation approach.

Yannig Goude: Adaptive probabilistic forecasting of net electricity demand

Electricity demand forecasting at day ahead horizon is essential to optimize production planning in accordance with future consumption. The development of (uncertain) renewable production and the increase of volatility of electricity markets induce a need for adaptive probabilistic net load (consumption minus renewable production) forecasting methods Hong and Fan [2016], Browell and Fasiolo [2021]. In this context, we focus on three significant aspects of electricity load forecasting. Firstly, our approach is adaptive, using models that incorporate the latest available observations. This enables our forecasting strategy to automatically adapt to changes in the system. Secondly, we prioritize probabilistic forecasting over point forecasting, recognizing the importance of uncertainty quantification in ensuring the efficient and reliable operation of electricity systems. Finally, we take into account both conventional load, which only considers consumption, and net-load, which takes into account embedded generation. Our approach builds on the success of the Kalman filter in adaptive point load forecasting De Vilmarest and Goude [2022], De Vilmarest and Wintenberger [2021], Obst et al. [2021]. To generate probabilistic forecasts, we use quantile regressions Fasiolo et al. [2021], Koenker and Bassett Jr [1978] on the residuals of the point forecasting model. Our methodology achieves adaptive quantile regressions using online gradient descent, where we utilize multiple learning rates and expert aggregation Wintenberger [2017] to avoid the need to choose a specific gradient step size. We apply this method to two datasets: the regional net-load in Great Britain and the demand of seven large cities in the United States Ruan et al. [2020]. Our adaptive procedures significantly improve forecast performance in both cases, for both point and probabilistic forecasting.

Thomas Muschinski: Robust weather-adaptive postprocessing using MOS random forests

Numerical weather predictions can have biases and miscalibrations that depend on the weather situation, which makes them difficult to effectively postprocess within the traditional model output statistics (MOS) framework based on parametric regression models. Consequently, much recent work has focused on the use of more flexible machine learning (ML) methods that are able to take into account additional weather-related predictors during postprocessing. Although these methods can achieve impressive predictive skill, they generally require significantly more training data than traditional MOS and are often less straightforward to implement and interpret. We propose MOS random forests, a new postprocessing method that avoids these problems by fusing the traditional and ML-based approaches. Rather than estimate properties of the forecast variable directly, as is commonly done, we use random forests to estimate the ”weather-adapted” coefficients of an assumed parametric base MOS. Since this base MOS contains valuable prior knowledge, much less data are required to effectively train the model and results are also easier to interpret. MOS forests are straightforward to implement and typically work well, even with no or very little hyperparameter tuning. The method is illustrated on the difficult task of forecasting daily precipitation sums in complex terrain.

Wednesday 10:30-12:00 – Machine Learning (chair Y. Goude)

Nicolas Lafon: A variational auto-encoder approach to sample multivariate extremes

Rapidly generating accurate extremes from an observational dataset is crucial when seeking to estimate risks associated with the occurrence of future extremes which could be larger than those already observed. Many applications ranging from the occurrence of natural disasters to financial crashes are involved. This paper details a variational auto-encoder (VAE) approach for sampling multivariate extremes. The proposed architecture is based on the extreme value theory (EVT) and more particularly on the notion of multivariate functions with regular variations. Experiments conducted on synthetic datasets as well as on a dataset of discharge measurements along Danube river network illustrate the relevance of our approach.

Moria Mayala: Subsampled base learners in extreme regions

In this work, we focus on the importance sampling, which allows us to overcome class imbalance issues for binary classification. Thus, as in Wager and al, we provide asymptotic normality properties in two samples cases using U-statistics and triangular arrays techniques.

Thibault Modeste: RKHS and some applications

In this talk, we will focus on the relations between RKHS and probability measures on Rd. We will be interested in kernels with a translation invariant Maximum Mean Discrepancy (MMD) which allows unbounded kernels. We will use the fact that these kernels k are not necessarily bounded to consider a stronger distance than usual, the Wasserstein distance W1.

Wednesday 17:00-18:30 – Environmental applications (chair C. Tillier)

Jonathan Koh: Predicting risks of temperature extremes using large-scale circulation patterns with r–Pareto processes

Many severe weather patterns in the mid-latitudes have been found to be connected to a particular atmospheric pattern known as blocking. This pattern obstructs the prevailing westerly large-scale atmospheric flow, changing flow anomalies in the vicinity of the blocking system to sustain weather conditions in the immediate region of its occurrence. Blockings' presence and characteristics are thus important for the development of temperature extremes, which are rarely isolated in space, so one must not just account for their occurrence probabilities and intensities but also their spatial dependencies when assessing their associated risk. Here we propose a methodology that does so by combining tools from the spatial extremes and machine learning literature, to incorporate 500hPa geopotential (Z500) anomalies over the North Atlantic and European region as covariates to predict surface temperature extremes. This involves fitting Generalized r-Pareto processes with appropriate risk functionals to high-impact positive and negative temperature anomaly events across central Europe from 1979–2020, using loss functions motivated by extreme-value theory in a popular boosting algorithm. We find which circulation patterns in the Euro-Atlantic sector are most important in determining the characteristics of these extremes, and show how they affect it.

Juliette Legrand: Bayesian spatio-temporal modelling of wildfire occurrences and sizes for climate model projections

Appropriate modelling of wildfire activity is critical in particular for risk managements. Here, we focus on wildfire risk in the Aquitaine region of France and its projection under climate change. More precisely, we aim to study whether wildfire risk could further increase under climate change in this specific region, which does not lie in the historical core area of wildfires in France, corresponding to the Southeast. For this purpose, we consider a marked spatiotemporal point process, a flexible model for occurrences and magnitudes of such environmental risks, where the magnitudes are defined as the burnt areas. The model is first calibrated on past observation data and then applied for projection of climate-change impacts using simulations of numerical climate models. We work within the framework of a spatiotemporal Bayesian hierarchical model, implemented using the INLA-SPDE approach. Results show a satisfactory fit of the model for the observation period. Depending on climate models, projections indicate low to moderate increase of wildfire activity under climate change. However, the increase is weaker than in the historical core area, which we attribute to different weather conditions (oceanic versus mediterranean).

Léo Pfitzner: Temperature forecasting with expert aggregations

A lot of Numerical Weather Prediction (NWP) models and their associated Model Output Statistics (MOS) are available. Expert aggregation (EA) has a lot of advantages to deal with all these models, like being online, adaptive to model changes and having theoretical guarantees. While EA works quite well in stationnary environments, it's quite difficult to deal with fast and large changes. We introduced a new expert aggregation algorithm - FSBOA - a combination of BOA (Wintenberger 2017) and FS (Herbster and Warmuth 1998), in collaboration with Olivier Wintenberger, Eric Adjakossa and Olivier Mestre. This new algorithm enables to have a second order regret bound in a non stationary environment. We also tested several aggregation strategies in order to improve the prediction of extreme temperature events like cold and heat waves. For this purpose, we added some biased experts of the Météo-France 35-member ensemble forecast (PEARP) to the set of models.

Thursday 10:00-12:00 – Forecast Evaluation (chair J. Broecker)

Sam Allen: Asymmetric proper scoring rules

Proper scoring rules provide an appealing framework with which to quantify the accuracy of probabilistic forecasts. Despite the wealth of literature on proper scoring rules, almost all scoring rules encountered in practice are symmetric. That is, if an event occurs with probability p, then a forecast that is smaller than p typically receives the same score (on average) as a forecast that is larger than p by the same amount. However, forecasters often wish to penalise under-prediction and over-prediction differently. For example, when forecasting high-impact weather events, a larger loss is typically incurred when the event occurs but is not predicted, rather than when the event is predicted to occur but does not. Although this asymmetry plays a key role in decision theory, it has received little attention in the context of probabilistic forecast verification, despite its obvious relevance. In this work, we introduce a general framework to construct asymmetric proper scoring rules for binary outcomes, which provides a simple, flexible, and intuitive generalisation of previously-proposed approaches. This framework is directly related to asymmetric problems in decision theory, and can be leveraged to obtain asymmetric versions of any proper scoring rule for binary outcomes. We further illustrate how this can be extended to form a general class of asymmetric scoring rules for real-valued continuous outcomes, which includes an asymmetric continuous ranked probability score (CRPS).

Harris Nkuiate: Evaluating Probabilistic Forecasts in the Presence of Observation Error

Observation error can cause proper scoring rules to be biased. Such scoring rules favour good forecasts of the observation rather than the truth, and yield scores that vary with the quality of the observations. Proper scoring rules can therefore favour worse forecasts of the truth and can reverse the ranking of competing forecasts in the presence of observation error.

Amaury Durand: Sequential evaluation of a forecaster by betting

We use a game-theoretic statistical framework to evaluate the performance of a sequential forecaster by betting against it. This is done by designing anytime valid sequential tests based on nonnegative supermartingales which represent our capital as we bet against the forecaster. By construction, such tests have time-uniform level guarantees meaning that, if the forecaster is good, there is a small probability that our capital will reach a certain threshold. More formally, the time at which our capital reaches the threshold (known as the rejecting time) is finite with small probability. Recent work provide asymptotic power guarantees which state that under certain alternative, the rejecting time will almost surely be finite. We improve such guarantees by providing an upper bound on the rejecting time under appropriate alternative.

Romain Pic: Distributional forecasting and its evaluation with the CRPS : Bounds and convergence of the minimax risk

The theoretical advances on the properties of scoring rules over the past decades have broadened the use of scoring rules in probabilistic forecasting. In meteorological forecasting, statistical postprocessing techniques are essential to improve the forecasts made by deterministic physical models. Numerous state-of-the-art statistical postprocessing techniques are based on distributional regression evaluated with the Continuous Ranked Probability Score (CRPS). However, theoretical properties of such evaluation with the CRPS have solely considered the unconditional framework (i.e. without covariates) and infinite sample sizes. We extend these results and study the rate of convergence in terms of CRPS of distributional regression methods. We find the optimal minimax rate of convergence for a given class of distributions and show that the k-nearest neighbor method and the kernel method reach this optimal minimax rate.