International Symposium of Forecasting 2025

June 29 – July 2, Beijing, China

General Chair: Yongmiao Hong
Local Chair: Feng Li
Program Chairs: Jue Wang and Yanfei Kang

Book of Abstracts

Group shrinkage for spatial autoregressive models with convex combinations of spatial weights matrices: an efficient Bayesian approach

Presenter: Xiaoyi Han (Xiamen University)
Author(s): Xiaoyi Han;Zhengzheng Cai;Jianchao Zhuo

Spatial autoregressive models with convex combinations of spatial weights matrices, including the dynamic spatial panels, have been employed to capture the relative importance of spillovers from different connectivity matrices in a variety of economic applications. In this paper, we investigate the scenario where different spatial weights matrices exhibit group structure, with each group representing a spatial spillover channel consisting of multiple connectivity matrices. We propose a new group shrinkage prior, the Group Inverse-Gamma Gamma with Bayesian Group Lasso prior, to identify the relevant versus irrelevant groups, pinpoint the influential connectivity matrices and mitigate multicollinearity within each group. We develop a computationally tractable Markov Chain Monte Carlo (MCMC) algorithm that integrates our new shrinkage prior with optimized hyperparameter tuning procedures, enabling accurate estimation and inference. We further propose an approximated exchange algorithm to enhance the computational efficiency. Simulation results demonstrate that our shrinkage prior outperforms existing priors in both estimation and variable selection, while the exchange algorithm proves computationally efficient. Empirical analysis examining channels of sovereign debt risk spillover also demonstrates the effectiveness of the proposed method.

Hierarchical Count Echo State Network Models with Application to Graduate Student Enrollments

Presenter: Qi Wang (University of California - Santa Cruz)
Author(s): Qi Wang;Paul Parker;Robert Lund

Poisson autoregressive count models have evolved into a time series staple for correlated count data. This paper proposes an alternative to Poisson autoregressions: count echo state networks. Echo state networks can be statistically analyzed in frequentist manners via optimizing penalized likelihoods, or in Bayesian manners via MCMC sampling. This paper develops Poisson echo state techniques for count data and applies them to a massive count data set containing the number of graduate students from 1,758 United States universities during the years 1972-2021 inclusive. Negative binomial models are also implemented to better handle overdispersion in the counts. Performance of the proposed models are compared via their forecasting performance as judged by several methods. In the end, a hierarchical negative binomial based echo state network is judged as the superior model.

An Infinite Hidden Markov Model with GARCH for Short-Term Interest Rates

Presenter: Qiao Yang (ShanghaiTech University)
Author(s): Qiao Yang;Chenxing Li

This paper introduces a novel Bayesian time series model that combines the nonparametric features of an infinite hidden Markov model with the volatility persistence captured by the GARCH framework, to effectively model and forecast short-term interest rates. When applied to US 3-month Treasury bill rates, the GARCH-IHMM reveals both structural and persistent changes in volatility, thereby enhancing the accuracy of density forecasts compared to existing benchmark models. Out-of-sample evaluations demonstrate the superior performance of our model in density forecasts and in capturing volatility dynamics due to its adaptivity to different macroeconomic environments.

Fault diagnosis of rotating machinery based on transfer learning and graph attention networks incorporating Kolmogorov-Arnold Networks

Presenter: Chen Yuan (Zhengzhou University)
Author(s): Yuan Chen

Graph theory has been widely applied to rotating machinery fault diagnosis in recent years. However, traditional graph models with multilayer perceptrons as the final feature extraction layer often suffer from fixed activation functions, limiting their flexibility in handling complex data patterns. Furthermore, available fault vibration signals for large rotating machinery are typically scarce. This paper addresses these challenges by proposing an innovative graph-based fault diagnosis model that integrates Kolmogorov-Arnold Networks and a transfer learning strategy. The proposed model features three key advancements: First, Visibility Graph technique is employed to transform fault vibration signals into network representations. Second, a hybrid graph attention network combining Kolmogorov-Arnold Networks with graph attention mechanisms captures both topological information and hierarchical deep features. Finally, transfer learning transfers knowledge from abundant source-domain data to target domains with scarce samples. Extensive experiments conducted on the CWRU, MFPT, and JNU bearing datasets validate the superior performance of the proposed model compared to other baseline models. The results emphasize the critical significance of integrating Kolmogorov-Arnold Networks with transfer learning strategies in graph-theoretic fault diagnosis frameworks.

Predicting Residual Fault Times in Reliability Growth Management under Limited Failure Data Set

Presenter: Wenjie Dong (Nanjing University of Aeronautics and Astronautics)
Author(s): Wenjie Dong;Kailing Li

Upon failure discovery, redesign or corrective measures are always implemented to eliminate the defects and improve system reliability in reliability growth management. In the presence of partial failure information or limited failure data set, however, it has been more and more challenging to confidently obtain accurate parameter estimators and reliability growth prediction results. Taking the S-shaped growth trend and saturation characteristics of reliability growth data into consideration, we investigate a two-stage reliability growth management procedure in predicting residual failure times and fault numbers. Specifically, Stage 1 is to fit the limited reliability characteristic quantity with the grey Verhulst model and to predict the subsequent reliability values within the range of system reliability requirement; whilst Stage 2 is to correspond the predicted reliability values with the Gompertz growth curve and to output the residual failure times of the reliability growth test. An illustrative example shows that integration of the Gompertz reliability growth curve and the grey Verhulst model is not only capable to effectively fit with the S-shaped failure data set, but also enables to precisely predict the residual failure times of the high-reliable product.

Interpretable Tree-based Cascade Ensemble Model for Breast Cancer Grading Prediction

Presenter: Huanze Zeng (Hohai University)
Author(s): Huanze Zeng;Binrong Wu;Haoyu Fang

Breast cancer is one of the most common malignant tumors among women, and early diagnosis and accurate grading are crucial for improving patient survival rates and optimizing treatment strategies. This study proposes an interpretable breast cancer grading prediction method based on a cascade ensemble of tree groups, aiming to enhance diagnostic accuracy and model interpretability. Initially, comprehensive relative importance analysis is used to select features from the input breast cancer imaging data, removing redundant or irrelevant features. Subsequently, these selected features are fed into a cascade ensemble of tree groups, which is composed of multi-layer classifiers, to enhance classification accuracy and robustness. Finally, through post-hoc feature importance ranking mechanism, we reveal the decision-making mechanism of the model, helping to identify the features that contribute most significantly to the predictions, thereby improving the model's transparency and credibility. Validation results show that this method demonstrates superior accuracy in breast cancer grading prediction compared to five other classical classification methods, providing significant support and reference for clinical practice.

Future Electricity Demand of Data Centers in South Korea: Based on the Basic Plan for Electricity Supply and Demand

Presenter: Cha-Ri Park (Korea Power eXchange)
Author(s): Cha-Ri Park;Ho-Seung Lee;Ah-Young Jeong

The rapid proliferation of AI technologies has led to a surge in demand for data centers, which is expected to have a significant impact on future electricity consumption. However, due to the recent and sudden increase in data centers, historical data is insufficient for accurately analyzing the growth trajectory of electricity demand. South Korea has separately incorporated the electricity demand of data centers in the 11th Basic Plan for Electricity Supply and Demand. This study provides an overview of South Korea's projected electricity demand for data centers in the Basic Plan for Electricity Supply and Demand. The analysis is based on historic electricity consumption data of data centers and the results of surveys on future data center construction intentions. To minimize overestimation in construction intention surveys, this study considers the feasibility of actual electricity supply by accounting for power system conditions. Additionally, assumptions regarding improvements in energy efficiency are incorporated through a literature review. This study provides a more realistic assessment of the rapidly evolving electricity demand of data centers and serves as a foundational resource for energy policy development and power infrastructure planning.

Long-Term Total Load Forecasting using a Functional Regression Model

Presenter: Haklim Shin (SKKU)
Author(s): Haklim Shin;Cha-ri Park;Ah-young Jeong

We propose a model for long-term forecasting of the daily total load curve, which represents the hourly total demand in South Korea. In this study, the daily load curve is assumed to be influenced by both deterministic factors–such as trends, temperature, and special days–and probabilistic factors, which account for time-dependent dynamic variations that cannot be explained by deterministic components. The deterministic factors are estimated on an hourly basis, reflecting their time-varying nature. Meanwhile, the dynamic variations in the total load curve, driven by probabilistic factors, are modeled using a functional autoregression estimation methodology. One-year rolling forecasting results demonstrate an average MAPE (Mean Absolute Percentage Error) of 2%, highlighting the robust performance of our proposed methodology.

Stochastic Modelling of Electricity Demand on Multiple Time Scales

Presenter: Zining Yuan (King's College London)
Author(s): Zining Yuan;Teemu Pennanen

This paper proposes a stochastic model for predicting half-hourly electricity demand 3 months to years ahead in Great Britain. The model that captures long-term trend, annual and daily seasonal patterns, as well as temperature dependencies in demand through semi-parametric additive models. It can incorporate expert views or other point forecasts to reduce forecasting errors. The model if easy to implement and computationally fast, which makes it

suitable for long term electricity price simulations and risk management. The model was developed in collaboration with the Low Carbon Contracts Company.

Fast Probabilistic Forecasting of Large Hierarchies

Presenter: Lorenzo Zambon (SUPSI, IDSIA)
Author(s): Lorenzo Zambon;Dario Azzimonti;Giorgio Corani

Often large hierarchies exhibit smooth upper-level time series but low-count, possibly intermittent bottom-level series, which are known to be difficult to forecast. Moreover, forecasting all the bottom series can be extremely time-consuming. We thus propose forecasting only a small subset of the upper series, where the signal is stronger, and propagating this information downward using a probabilistic top-down approach. This approach produces a joint bottom-level forecast distribution in the form of samples, using in-sample empirical distributions to disaggregate the upper-level forecasts. Since the bottom samples are joint, probabilistic forecasts for any other upper level can then be easily obtained. We validate our method on the M5 dataset, which includes approximately 30,000 bottom-level time series, over 90% of which are lumpy or intermittent (ADI > 1.32). We forecast only a few hundred series using ETS, without using explanatory variables. We then apply our probabilistic top-down algorithm to generate 2,000 coherent samples from the joint forecast distribution; this requires only a few minutes on a standard laptop. Our method achieves an excellent performance, ranking among the top 20 in the M5 uncertainty competition.

Devaraj-Fitzgerald Intermittent Demand Forecast Method: An adaptive, statistical approach to accurate intermittent demand forecasting

Presenter: Renold Raj Devaraj (Dell Technologies)
Author(s): Renold Raj Devaraj;Jeannie Fitzgerald

Demand forecasting is a critical component of supply chain planning, particularly for industries handling low-volume, high-cost products, where forecasting errors can lead to significant financial and operational repercussions. The Devaraj-Fitzgerald Forecast Method addresses these challenges by leveraging foundational statistical principles to develop a non-complex, interpretable, and computationally efficient forecasting model. The approach begins with a descriptive statistical segmentation of the time series data using an empirically derived threshold to distinguish between periods of demand and low or zero demand. For each segment, the method dynamically identifies best-fitting probability distributions using a suite of robust tests, including Kolmogorov-Smirnov, Anderson-Darling, and Cramer-Von-Mises. Probability distributions considered include, but are not limited to, gamma, Weibull, log-normal, normal, and logistic Forecasts are generated by computing a weighted average of the top ranked probable distributions for each segment, thereby capturing the nuances of both regular and irregular demand patterns. These forecasts are then aligned with future time indices using a non-linear regression model, ensuring continuity and adaptability across various temporal scales (daily, weekly, or monthly). Like traditional intermittent forecasting techniques like SBA, Croston, and SBJ, this method reduces computational complexity while maintaining accuracy (we report an average improvement of 6-10% when compared to SBA/Croston and SBJ methods), making it an ideal complementary tool for a diverse forecasting toolkit. By accurately modelling demand for irregular time series data, the Devaraj-Fitzgerald Forecast Method enhances inventory planning, optimizes production schedules, and ultimately contributes to improved customer satisfaction and reduced operational risk.

Evaluating and boosting the reliability of global neural forecasting models

Presenter: Jente Van Belle (KU Leuven)
Author(s): Jente Van Belle;Wouter Verbeke;Nikolaos Kourentzes

Global forecasting models are trained by minimizing forecast errors across all time series in a given dataset. The exact weighting of individual series depends on time series preprocessing, sampling procedures, and other modeling considerations; however, it is typically not explicitly designed by the modeler. We argue that this lack of control may lead to reliability concerns, as the model may distribute focus unevenly across series, and large errors may also go unnoticed. Of particular relevance here is the heterogeneity in the training dataset and the presence

of subgroups of series within it. Avoiding large errors and subgroup underperformance is crucial, as these can significantly impact downstream operations and erode confidence in the forecasting system. Moreover, as operational decision-making becomes increasingly automated, these forecast model reliability concerns should be treated as a key consideration rather than an afterthought. In principle, these concerns do not apply to local forecasting models, as they are built independently for each time series. However, in operationalizing local models, similar issues can arise, as unreliable forecasts may be masked by summary statistics used for performance evaluation. In this work, we focus on developing methodologies to measure and improve the reliability of global forecasting models. Specifically, we propose (i) a systematic data-centric procedure to evaluate model reliability and (ii) methodologies that integrate forecast model reliability directly into the training process. The developed data-centric evaluation framework—which focuses on subgroups of series—may nonetheless also be useful for more transparent performance summaries for local models.

Forecasting company cash flow in the logistics industry utilizing machine learning techniques

Presenter: Alice Wolfe (IESEG School of Management, Univ. Lille, CNRS, UMR 9221 – LEM – Lille Economie Management, F-59000 Lille, France)
Author(s): Alice Wolfe;Sarah Van der Auweraer;Kristof Coussement

Cash flow management of a company is a critical operation that impacts daily operations, spending, investments, revenue and risk management. With enhanced cash flow forecasting companies can improve their cash management, becoming more agile in their investment decisions and operations, which can bring significant increases in revenue. Despite its high importance to any operating company, it remains a limited field of study within academic forecasting research and has largely focused on the utilization of traditional forecasting methods. Cash flow forecasting lags behind general financial forecasting in the exploration of Machine Learning methods and their potential improvement to the forecast. The aim of this research is therefore to inspect the application of Machine Learning methods on company cashflow forecasting as compared to traditional forecasting methods. This research is carried out in close collaboration with an international logistics company in the fashion and lifestyle industry. The company provided a data set of approximately 3 million financial records comprised of client and supplier invoice and payment records from 2016 to 2025. This context-specific data is combined with factors regarding contextual client and company data. Using this data, we test and compare traditional forecasting methods: ARIMA and ETS, with Machine Learning methods: LSTM, SVR, and Gradient Boosting, to provide a forecast of the company's cash flow across different time horizons.

Machine learning in foreign exchange

Presenter: Bo Yuan (University of Cambridge)
Author(s): Bo Yuan

This paper investigates the application of machine learning (ML) techniques to predict the cross-section of currency excess returns and to better understand the sources of FX risk premia. We address three main questions: (1) Do ML models statistically improve upon the predictive performance of traditional linear models? (2) Can ML-based investment strategies generate economically meaningful gains in risk-adjusted returns? (3) Can interpretability tools uncover the key factors and interactions that drive return variation? We employ neural networks and other machine learning algorithms to model currency excess returns using a large set of predictors that include both tradable factors (e.g., carry, momentum, value) and nontradable macro-financial indicators (e.g., interest rate differentials, inflation surprises, global risk measures). Neural networks outperform all benchmark models, including linear regression and random forests, across most forecast horizons. Importantly, they are the only models that consistently beat the random walk benchmark in out-of-sample prediction, particularly at short horizons. Predictive power diminishes with longer horizons, aligning with the declining signal strength observed in international finance literature. Our results extend prior work by incorporating a broader universe of 49 currencies and a richer set of predictors than previous studies. We also demonstrate the feasibility of implementing ML-based currency portfolios, which deliver higher Sharpe ratios than conventional factor-based strategies. Our findings suggest that machine learning not only improves prediction in FX markets but also sheds light on the underlying economic mechanisms—providing a pathway to reconcile statistical accuracy with economic interpretability. To address concerns about model transparency, we apply both local (e.g.,

DeepLIFT, Layer-wise Relevance Propagation) and global (e.g., Shapley value) interpretability methods. These tools help identify which variables contribute most to predicted returns, and how the interaction between global macro conditions and currency-specific characteristics shapes forecast accuracy. We find that no single predictor dominates across currencies or horizons; rather, predictive performance hinges on nonlinear interactions between groups of factors. For instance, the effect of carry depends critically on global liquidity and risk sentiment.

Optimal Pricing during Sales-Events in the Online Fashion Industry

Presenter: Tim Januschowski (Databricks)
Author(s): Adele Gouttes; Mones Raslan; Stefan Birr; Manuel Kunz; Tim Januschowski

We present a novel approach to pricing optimization for e-commerce sales events at Zalando.Our submission describes our transition from a traditional pricing system based on analytics,heuristics, and manual adjustments to a fully automated algorithm that explicitly balancesrevenue objectives with profitability. Our solution addresses the unique challenges of salesevents, which require more flexible and responsive pricing than regular periods due to theirirregular length, sudden demand changes, and higher frequency of steering interventions.The proposed system consists of two main components:(1) a high-frequency sales forecasting model using Gradient-Boosted Regression Trees thatpredicts daily sales at the article level across different discount levels, and(2) the algorithm that optimizes discount selections by balancing short-term revenue withlong-term profitability.Our forecasting model is adopted to work optimally in this forecast-than-optimise setting andtuned to improve accuracy, particularly for articles with sparse sales. The optimization algorithmuses these forecasts to determine Pareto-efficient discount recommendations that maximize aweighted combination of Net Merchandise Value and long-term profit.Through extensive A/B testing across multiple markets and sales events, we demonstrate thatour approach achieves a 6% increase in per-sales-event profit while maintaining or slightlyincreasing revenue compared to the previous system. The algorithm's flexibility allows it to runin minutes rather than hours, enabling rapid responses to changing market conditions. Ourimplementation successfully addresses the challenges of the fast-paced, high-stakes salesevent environment while providing transparent, evidence-based pricing decisions.This work contributes to the field of algorithmic pricing by demonstrating how machine learningforecasts can be effectively combined with optimization techniques to balance multiple businessobjectives in a complex retail environment with millions of time series and frequent cold-startproblems.

A Dynamic Knowledge-Enhanced Data-Driven Model For Marine Oil Spill Trajectory Forecasting

Presenter: Liu Wenzhi (China University of Petroleum)
Author(s): Wenzhi Liu;Xiaofeng Xu;Lean Yu;Kai Zhang;Zhiting Liu;Liming Zhang

Marine oil spill trajectory is difficult to forecast accurately due to the superimposition of various physical, chemical and biological processes, while the complex linear and nonlinear features exhibited in trajectory data. Consequently, traditional methods struggle to the accuracy dilemma between global and local trajectory forecasting, which makes it difficult to realize large-scale accurate forecasting. To solve this problem, this paper proposes a novel dynamic knowledge-enhanced data-driven (DD) forecasting framework. Compared with traditional model-averaging, the designed framework is not a static combination of the two models' final results, but a dynamic fusion throughout the forecasting process. That is, the knowledge-driven (KD) model performs independent oil spill forecasting with each DD result as an initial state and corrects it. In addition, in order to improve the independent forecasting performance of each model, this paper designs a data enhancement strategy based on case-based reasoning (CBR) as well as a 3D simulation scheme for the oil spill motion process. The former improves the forecasting accuracy and stability of the DD model by clustering the training data as well as dynamically updating it. The latter integrates the processes of horizontal diffusion and vertical transport to improve the forecasting accuracy and realism of the KD model. Finally, to address the differences in the output data forms of the two models, this paper unifies the two by using the mass transformation formula and bilinear interpolation method. To evaluate the performance of the proposed method, real oil spill data from the Deepwater Horizon incident was used. Experimental results demonstrate the superior performance of the developed method compared to existing method. Furthermore, the dynamic fusion method significantly improves forecasting accuracy and robustness over a wide range.

Shaping Honesty: Strategies for Truthful Reporting in B2B Electricity Saving Platsforms

Presenter: Yang Seongwon (SKKU)
Author(s): Seongwon Yang;Yeji Lim;Kyungsik Nam;Cha Ri Park;Ah Young Jeong

We model the demand-response program in an incomplete information game between the utility and demand response aggregators who represent consumers. The source of incomplete information originates from the unobservable demand and accordingly unobservable demand reduction during the event periods. We solve the optimal baseline and rebate that minimize the procurement cost under asymmetric information. Our optimal baseline and rebate solution incentivize demand response aggregators and their consumers to reveal their actual demand and reduce their moral hazard of inflating the reduction to increase the rebate payments. Using the empirical data for our theoretical solution, we estimate demand response aggregator's utility function and planner's cost function, and measure improved cost-efficiency relative to the case of using the observed baseline. The policy implication is also provided for efficient demand response systems.

A CoT-Empowered LLM for Crude Oil Return Prediction: Evidence from China

Presenter: Hanfeng Zhou (AMSS Center for Forecasting Science, Chinese Academy of Sciences, Beijing, 100190, China)
Author(s): Hanfeng Zhou;Yuting Wei;Han Feng;Pengyang Song;Yinuo Ren;Jue Wang

Chain-of-Thought (CoT) prompting has emerged as a powerful technique for enhancing the reasoning capabilities and interpretability of large language model (LLM), particularly in complex zero-shot tasks. In this study, we present a CoT-augmented large language model (LLM) framework to extract predictive signals from unstructured news texts for crude oil return forecasting. First, we employ LLMs guided by CoT prompting to identify causal linkages between discrete news items and crude oil price dynamics—laying the analytical groundwork for downstream evaluation; second, building on these causal insights, we develop a multi-dimensional factor quantification mechanism, where each news item is assessed along key economic dimensions using domain-informed prompts, yielding interpretable and high-resolution numerical descriptors; third, we conduct temporal aggregation to synthesize daily factor representations, enabling their seamless integration with traditional time-series predictors. Experimental results demonstrate that incorporating domain-specific energy economics knowledge through CoT reasoning significantly enhances zero-shot extraction of predictive factors, with consistent and substantial improvements in out-of-sample predictive performance across multiple forecast horizons—daily, 2-day, and weekly (5-day) forecasts. Further analyses examine the information overlap between news-derived factors and traditional numerical data. Additionally, we perform perturbation experiments to assess the impact of sentiment injection in prompts and evaluate the framework's stability when exposed to anomalous domain knowledge.

Chronological Valuation Framework for Artwork Property Rights

Presenter: Yuan ni (Beijing Information Science and Technology University)
Author(s): Yuan Ni;Rui Wang;Jian Zhang;Peng Zhou;Xiao-Min Lv

The evaluation of the property rights value of paintings is of great significance for promoting the prosperity of the cultural art market and ensuring the orderly transfer of cultural assets. Aiming at the problems that the existing evaluation models for the property rights value of paintings have insufficient mining of sparse data features and are difficult to capture the multi-dimensional dynamic changes in the value of paintings, this paper proposes a Multimodal TimeGAN Graph Convolutional Neural Network Model (MTG-GCN). Firstly, based on the multi-source heterogeneous features of paintings, a dynamic similarity retrieval model is constructed to generate a set of prices of similar paintings in the time dimension. Secondly, the Time Series Generative Adversarial Network (TimeGAN) is used to build a simulation time series generation model for paintings, which transforms the set of similar paintings at discrete time points into continuous time series simulation data. Finally, based on the simulation time series data, a painting property rights value evaluation model based on the time series graph convolutional neural network is constructed to achieve the dynamic value evaluation of painting property rights. The experimental results show that compared with traditional methods, this model performs excellently in the retrieval accuracy of similar paintings and

the preservation of time series features, and its evaluation results have a significant correlation with the actual market value trends.

Evolution and response of flood disaster network public opinion based on an event logic graph: A case study of rainstorms in Henan, China

Presenter: Li Cui (University of Emergency Management)
Author(s): Cui Li

Social media takes the center stage as the primary information and opinion-sharing platform for the public during disaster events. Related topics quickly spread, and multiple public opinion events exhibit clear coupling phenomena, leading to the formation of network public opinion. By creating a network public opinion event logic graph, our study illustrates the development and evolution of network public opinion during disaster events and provides the government with a crucial foundation for action. For this study, we selected tweets on Sina Weibo regarding the “Henan rainstorm” between July 19, 2021, and August 2, 2021. To construct a network public opinion event logic graph, pattern matching was used to discover event relationships and extract causal event pairs and sequential event pairs. Moreover, the term frequency-inverse document frequency + K-means clustering algorithm (TF-IDF + K-Means) clustering algorithm was used to group event pairs with high similarity into one category. The critical nodes and evolution pathways in the event logic graph were then thoroughly examined. Our findings demonstrated that the topic evolution of network public opinion was multidirectional in nature, covering topics such as disaster events, emergency rescue, public sentiment, media opinion, and livelihoodrelated events. The two main components of the evolutionary pathway of network public opinion are the evolution of the disaster itself and its evolution in the social sphere. Our study provides a technique for determining the dynamic evolution of network public opinion and provides useful consultation for the government’s upcoming response to network public opinion during disaster events.

Research on Pricing Models of Data Trading Platforms Considering Business Model Differences: A Two-Stage Dynamic Game Modeling Approach

Presenter: Yan Zhao (Beijing Technology and Business University)
Author(s): Yan Zhao;Yue Ma

In light of the differences in user ownership and profit models under distinct business models of data trading platforms, this study examines three pricing strategies—registration fees, transaction fees, and a hybrid fee structure—and constructs a two-stage dynamic game pricing model based on the theory of two-sided markets. Numerical simulations are employed to analyze the optimal pricing processes for both information intermediary platforms and one-stop service platforms. The findings indicate that: (1) Under the information intermediary business model, data trading platforms should adopt a hybrid pricing strategy, with service process optimization and cost reduction serving as the primary pathways to enhance platform profitability; (2) Under the one-stop service business model, data trading platforms can achieve higher profits by integrating hybrid pricing with data mining, data matching, and other data-driven information technologies; (3) Traditional information intermediary platforms can transition to one-stop service platforms by implementing a hybrid pricing model, and this business model innovation can break through the profit limitations of conventional data trading platforms.

Conformal-Prediction-based Probabilistic Forecasting for Time Series: Performance Evaluation of Novel Models Against Bootstrapping and DeepAR

Presenter: Pedro Leal (Departamento Administrativo Nacional de Estadística - DANE Colombia)
Author(s): Pedro Leal;Mario Arrieta

This paper introduces CPnewModel and CPnewModelQn, two probabilistic forecasting frameworks that integrate conformal prediction with time series decomposition and adaptive regularization. While CPnewModel leverages bootstrapped residuals for uncertainty quantification, CPnewModelQn replaces resampling with a sorted empirical Qn distribution, prioritizing computational efficiency. We evaluate both models against Bootstrap and DeepAR

across 120 simulated ARMA scenarios with non-Gaussian noise (uniform, exponential, t-distributed, and mixtures). Results show CPnewModelQn excels in nonparametric settings, outperforming DeepAR in accuracy (ECRPS: 0.91 vs. 1.26, $p < 0.001$) and stability (std. dev.: 0.48 vs. 1.90). Statistical testing via Wilcoxon and Nemenyi procedures confirms CPnewModelQn's robustness to distributional shifts, bridging classical time series analysis with modern machine learning. These advances highlight the potential of conformal methods to deliver sharp, adaptable forecasts in complex temporal environments.

Non-Stationary Capacity Forecasting for Cloud Storage Resource Pools Based on Probability Characteristic Modeling of Stochastic Jumps

Presenter: Wang Shi (lenovo)
Author(s): Wang Shi;Ming Lu;Zheyi Zhu

Cloud capacity planning relies on high-quality forecasting methods for corresponding architectural design, equipment selection, operational strategy formulation and reliability planning. Affected by factors such as sudden capacity expansion and resource delivery or migration, the capacity of cloud storage resource pools shows non-stationary characteristics and is difficult to forecast accurately. This research proposes a non-stationary capacity forecasting method for cloud storage resource pools based on probability characteristic modeling of stochastic jumps. The method initially constructs a multi-stage linear regression based on the lifecycle of the resource pool to depict the growth trend of capacities across different lifecycles. It then models the probability characteristics of stochastic jump time intervals and the number of jumps using the Compound Poisson Process, combining jump variables to build a renewal reward process, thereby achieving the modeling of non-stationary stochastic jump capacities. Furthermore, it derives the distribution expression of cumulative jump capacities based on the modified Bessel function and obtains capacity forecasting results at different probability levels through statistical simulation. Compared to traditional deterministic forecasting methods, this method achieves effective forecasting of non-stationary capacities and accurately captures capacity changes in storage resource pools under different lifecycles and unexpected maintenance actions. By continuously collecting relevant monitoring data from dozens of proprietary data centers and cloud platforms worldwide over several years, this method has been verified to accurately forecast the first arrival time to the threshold and future distribution of storage capacities. It also provides probabilistic confidence interval data to assist cloud computing operation engineers in capacity planning for storage resource pools, thereby enhancing the reliability level of storage services. At the conclusion of the report, we will present cases of application of this technology in capacity planning, monitoring, and alerting scenarios within cloud computing intelligent operations and maintenance.

Prediction in a Double-Nested Error Components Panel Data Model

Presenter: None
Author(s): Alain Pirotte

Prediction is often an integral part of an empirical work in economics. Nevertheless, prediction in the context of multi-dimensional panels, especially hierarchical ones, remains unusual. In the spirit of Baltagi and Pirotte (2013), this paper studies the properties of the Best Linear Unbiased Predictor (BLUP) for a double-nested error components panel data model employing Monte Carlo simulations. Having an overview of these properties is important because it is now common to use panel data characterizes by a multi-dimensional nested structure. Examples include predicting average price of a box of 2kg corresponding to a product in a chemical class, or economic growth of a NUTS3 region in a NUTS2 region in a NUTS1 region in Europe. Using Monte Carlo simulations, we show that this predictor performs better in terms of root mean square error performance than the usual fixed and random effects predictors ignoring the nested structure of the data, and the one associated with a single-nested error components structure.

Forecasting the age structure of the scientific workforce in Australia

Presenter: Rob Hyndman (Monash University)
Author(s): Rob Hyndman;Kelly Nguyen

I will demonstrate how to forecast the age structure of a workforce using demographic growth-balance equa-

tions that incorporate functional data components for graduate influx, international mobility, retirement, and mortality. This model adapts the stochastic population framework proposed by Hyndman and Booth (IJF, 2008), substituting the birth component at age 0 with a graduation process at any age, and integrating retirement dynamics alongside a standard mortality process. It accommodates variable graduate influx, international mobility, and aging, through functional data components. The model is illustrated by forecasting the future age structure of scientific disciplines in Australia, utilising a variety of data sources. Data limitations are addressed through cohort interpolation and constrained penalised regression splines in order to obtain a workable data set. A combination of local and global ARIMA models is employed to address the temporal dynamics in the time series components. By forecasting age structures in disciplines such as Physics and Astronomy, Mathematical Sciences, Chemical Sciences, Earth Sciences, and Biological Sciences, the model informs strategic workforce planning and policy development. This forecasting framework not only aids in strategic planning and policy formulation within the scientific community, but also offers a blueprint for other professional sectors seeking to understand workforce dynamics amid demographic shifts.

Mortality Forecasts to Verify Recovery in the Life Insured Mexican Population

Presenter: Kristian Frich (Universidad Anáhuac Norte Student)
Author(s): Kristian Frich;Eliud Silva;Fernando Sánchez

The impact of the COVID-19 pandemic on life expectancy has been one of the most critical challenges faced by modern society. Like many other countries, Mexico experienced the consequences of the health crisis, reflected in a decline in life expectancy, both insured and non-insured population. The lack of comprehensive studies in our country for life insured populations motivated us to estimate and forecast potential recovery levels in life expectancy following the pandemic. The data were taken from the Mexican Association of Insurance Institutions (AMIS). They covered the period from 2014 to 2022, including exposure and mortality records for individual and group life insurance policies, broken down by gender. Considering data before the pandemic, we aim to model and forecast life expectancy in a hypothetical scenario without the COVID-19 pandemic. For that, the Lee-Carter model was employed, so we forecast from 2020 to 2025. Our results highlight that pre-pandemic mortality levels have not yet been recovered overall. This was identified by comparing the observed life expectancy with the forecasted values under the mentioned scenarios. In particular, the gap between actual and expected life expectancy remains noticeable for both individual and group life insurance policies, suggesting that the effects of the pandemic continue to impact insured populations beyond the initial health crisis.

Using Google Trends Data to Forecast Homicide Mortality: The Case of Mexico

Presenter: Eliud Silva (Universidad Anahuac México)
Author(s): Eliud Silva;Eduardo Vázquez

This study aims to compare traditional forecasting models —ARIMA and Vector Autoregressive (VAR)—with and without Google Trends data, the research explores ways to enhance prediction accuracy. Using homicide records from INEGI and Google Trends data from 2006–2018, the study highlights the integration of real-time online data to complement official statistics. Considering a forecast horizon of 15 months, results show that while ARIMA models benefit slightly from Google Trends, VAR models significantly improve, especially for female homicide forecasts, reducing errors substantially. For male homicides, ARIMA models with Google Trends perform best. The findings underscore the potential of integrating digital data sources into traditional models to provide more accurate and timely tools for public safety planning and intervention.

Bayesian quantile inference and order shrinkage for hysteretic quantile autoregressive models

Presenter: Kai Yang (Changchun University of Technology)
Author(s): Kai Yang;Bo Peng;Xiaogang Dong;Chunjing Li

Hysteretic quantile autoregressive model combines the hysteretic patterns and quantile autoregression, which can capture the dynamic relationship and nonlinear characteristics at different quantiles in time series data. In this

paper, the Bayesian quantile inference and order shrinkage are studied for a class of hysteretic quantile autoregressive time series models. By using Markov Chain Monte Carlo (MCMC) techniques, the proposed Bayesian quantile method can handle the sparse hysteretic quantile autoregressive model well. It can accurately determine order of the model and estimate nonzero coefficients. Both simulation studies and a data example show that the proposed methods are feasible, reliable and appropriate for analysing the US Gross National Product data set.

Quantile regression for count time series forecasting

Presenter: Sheng Danshu (Harbin Institute of Technology)
Author(s): Danshu Sheng

Estimating conditional quantiles plays a crucial role in modern risk management and other various applications. However, the quantile regression (QR) estimation of Poisson autoregressive (PAR) models, count-type models, remain an unresolved challenge. In this study, we propose a novel approach that employs a jittering smoothing method and a novel transformation strategy to convert this complex problem into an easily implementable quantile regression problem for continuous-type regression models. The asymptotic theory of the estimator is derived under some regularity conditions and the applications to four popular and classical PAR models are considered. Additionally, a novel h-step prediction method (h-QRF) is developed to forecast the h-step conditional distribution. The finite sample performance of the method is examined, and its advantages over existing methods are illustrated by simulation studies and an empirical application to the daily stock volume dataset of Technofirst.

Bayesian Statistical Inference for Spatial Panel Varying-Coefficient Quantile Regression Model

Presenter: Zhuoxi Yu (Liaoning University)
Author(s): Zhuoxi Yu;Shuang Yang

We study the spatial panel varying-coefficient quantile regression model under the assumption that the error term follows an asymmetric Laplace distribution in this paper. By using B-spline functions to approximate the varying-coefficient functions, the MCMC algorithm combining Gibbs and Metropolis-Hastings sampling is applied to obtain the Bayesian nonparametric estimation of the varying-coefficient functions. Numerical simulations demonstrate the effectiveness of the nonparametric estimators. Furthermore, this method is applied to analyze the relationship between regional GDP and the three major industrial structures. The empirical results also confirm that the proposed method is both feasible and reliable.

Sparse Heterogeneous Auto-Regressive Model for Volatility Forecasting

Presenter: Mingmian Cheng (Sun Yat-sen University)
Author(s): Mingmian Cheng

This paper introduces modifications to improve the performance of the widely-used Heterogeneous Auto-Regressive (HAR) model in volatility forecasting. While our approach, similar to the classic HAR model, incorporates volatility components aggregated over different time horizons, it results in a sparse AR-type representation for volatility, in contrast to the full AR-type structure implied by the classic HAR model. Consequently, we refer to it as the Sparse HAR (SHAR) model. The SHAR model is implemented through a three-step procedure. First, significant lag terms within an AR model of daily realized volatility are identified using shrinkage methods. Next, these selected lags are combined across various time horizons. Finally, predictors with significant power are extracted from the aggregated volatility components produced in the previous step. To mitigate the inherent instability in variable selection throughout the implementation procedure, stability selection is employed in the initial step and bootstrap aggregation in the final step. Empirical results demonstrate that our SHAR model significantly improves predictive accuracy, particularly for medium- and long-term forecasts, with pronounced benefits observed during periods of prevailing financial market turmoil.

Heterogeneous Predictability on Mutual Fund Alphas: A Sparse Clustering GMM Approach

Presenter: Jiangshan Yang (City University of Hong Kong)
Author(s): Jiangshan Yang;Liyuan Cui;Guanhao Feng

Mutual fund managers' skills, represented by alphas for risk-adjusted performance, are predictable, time-varying, and influenced by various fund characteristics identified via machine learning (e.g.,kaniel2023machine,demiguel2023machine). Recent literature often overlooks the heterogeneous data structure in the cross section of mutual fund alphas, as most funds lack positive alphas. Additionally, the ability of fund characteristics to predict alphas may vary over time between skilled and unskilled managers. This paper introduces a novel nonparametric clustering method to identify and estimate latent group structures, effectively distinguishing a small set of skilled fund managers from their counterparts. Fund managers are grouped based on their market timing skills, allowing those within the same group to share group-specific parameters related to aggregate market predictors. Overall, our approach, Sparse Clustering GMM (SCGMM), facilitates the simultaneous identification and estimation of group-specific parameters, heterogeneous functional forms, and sparsity in a potentially high-dimensional set of covariates, all without the need to pre-specify functional forms or the number of groups. Finally, our empirical study using U.S. mutual fund data reveals that only a small proportion of funds possess predictable alphas, dependent on a limited number of aggregate market predictors.

Multi-Horizon Test for Market Efficiency

Presenter: Xiye Yang (Rutgers University)
Author(s): Xiye Yang;Merrick Li

In efficient markets, asset returns show no predictability over short periods, such as daily or intraday intervals, where expected returns are nearly zero. However, the presence of pricing errors—transitory components of asset prices reflecting various market frictions or biases—induces return reversals. These reversals provide evidence against market efficiency, particularly in terms of liquidity provision. We propose to test market efficiency based on the joint inference of the covariances of observed returns over multiple horizons. We show that analyzing a small set of horizons is sufficient to detect inefficiencies caused by a wide range of transitory pricing errors, both theoretically and practically. Extensive simulations highlight the superiority of our multi-horizon approach over traditional tests, especially when pricing errors are weak and exhibit complex serial dependencies. Moreover, our test statistic serves as a natural liquidity measure that can effectively identify financial crisis with significant liquidity drains.

The Sign vs. Magnitude Nexus in Assessing Predictability

Presenter: Constantin Burgi (University College Dublin)
Author(s): Constantin Burgi;Dimitris Thomakos

The evaluation of forecasting performance rests on the use of both descriptive measures and formal inference of a model/method vis-a-vis a benchmark. Standard practice tends to look at measures of sign accuracy like market timing and magnitude accuracy like mean absolute or squared errors separately, but these two components are naturally linked. In the context of economic time series such as inflation, or financial time series as for example the returns of a tradable asset, forecasting signs accurately is a prerequisite for having any assessment on magnitudes. Thus, forecast evaluation based on an approach that combines signs and magnitudes might offer better insights about model/method performance than a standard evaluation based on only one of the two. In this paper we offer a coherent methodology of addressing this idea, by providing a number of test procedures that assess magnitude accuracy conditional on sign accuracy and compare the magnitude performance on the observations whose signs have been correctly predicted vs. ones whose signs where incorrectly predicted. The interpretability of our suggested test procedures is illustrated with simulations and the empirical applicability is illustrated with a number of real-world series where sign accuracy appears to be a prerequisite for magnitude accuracy. Our results suggest that their use can lead to more robust and meaningful model/method rankings in forecasting comparisons.

Partial Time-Varying Regression Modelling under General Heterogeneity

Presenter: Yufei Li (King's College London)
Author(s): Yufei Li;Liudas Giraitis;George Kapetanios;Tien Chuong Nguyen

This paper explores a semiparametric version of a time-varying regression, where a subset of the regressors have a fixed coefficient and the rest a time-varying one. We provide an estimation method and establish associated theoretical properties of the estimates and standard errors in extended for heterogeneity regression space. In particular, we show that the estimator of the fixed regression coefficient preserves the parametric rate of convergence, and that, de- spite of general heterogenous environment, the asymptotic normality property for components of regression parameters can be established and the estimators of standard errors have the same form as those given by White (1980). The theoretical properties of the estimator and good finite sample performance are confirmed by Monte Carlo experiments and illustrated by an empirical example on forecasting.

Mixed Membership Estimation in Partial Correlation Network

Presenter: Siao Xu (University of Mannheim)
Author(s): Siao Xu

We propose a two-path partial correlation model (2PPCM) to learn the mixed membership structure in the panel. The model is assumed to follow a joint distribution specified by a Bayesian belief network. Correlation structure and interaction intensity are determined by two paths independently. We put forward a PartiaCorr-mixed-SCORE Algorithm to estimate the mixed memberships with 2PPCM. The consistency of the algorithm is also established and the simulation proves its effectiveness. We learn the industry mixed membership structure of all U.S. states using U.S. employment growth data.

On the Estimation of Forecaster Loss Functions Using Density Forecasts

Presenter: Wuwei Wang (Southwestern University of Finance and Economics)
Author(s): Wuwei Wang;Kajal Lahiri;Fushang Liu

We suggest a novel approach to use density forecasts from surveys to identify asymmetry in forecaster loss functions. We show that we can calculate the loss function parameters for Lin-Lin and Quad-Quad loss functions based on the first order condition of forecast optimality. Since forecasters form their point forecasts based on what they believe to be the data generating processes and their loss functions, we can reverse this process and learn about forecaster loss functions by comparing their point forecasts and density forecasts for the same target. The advantage of this method is that we can relax the two assumptions needed in Elliott, Timmermann and Komunjer (2005) and Elliott, Komunjer and Timmermann's (2008) GMM method: the point forecasts and density forecasts need not to be rational and the loss function parameters need not to be constant over time. Moreover, we do not need to know the actual values of the target variable. This method is applied to density forecasts for annual real output growth and inflation obtained from the Survey of Professional Forecasters (SPF) during 1968-2023. We find that forecasters treat underprediction of real output growth more dearly than overprediction, reverse is true for inflation.

Testing Purchasing Power Parity using a general non-linear model

Presenter: Bo Guan (Cardiff Business School)
Author(s): Bo Guan;Saeed Heravi;Khelifa Mazouz;Joshy Easaw

This study proposes a general non-linear State Dependent Model (SDM) as an alternative method of testing the Purchasing Power Parity (PPP). We show that the real exchange rate adjustment process exhibits a mean reversion behaviour for Pound, Euro and Yen, while Brazilian Real follows a random walk. We also use a non-linear SDM model to examine the behaviour of these exchange rates under different levels of global uncertainties. The results suggest that the speed of adjustment to PPP is an increasing function of global uncertainty.

Selection of Tourism Demand Analysis Following a Prolonged Economic Crisis

Presenter: Saeed Heravi (Cardiff Business School, Cardiff University)
Author(s): Saeed Heravi;Chen Xu;Bo Guan;Peter Morgan;Antonio Almeida;Brian Garrod

ABSTRACT Covid-19 has introduced greater non-linearity and multiple structural breaks into datasets, making them less serviceable for forecasting. At the regional level, the situation is further confounded by poor data quality. This paper aims to demonstrate the potential of the state-dependent model (SDM) to produce robust international and regional tourism demand forecasts following prolonged economic crises. We describe how the state-dependent model can be used in prediction, and the predictive performance of the SDM is first tested on simulated time series data. The SDM is then assessed using highly seasonal data on inbound tourism to China covering the SARS crisis of 2002 to 2004 and tourism arrival data to Madeira Island showing very little seasonal pattern and covering the Covid-19 global pandemic. Forecasts are then developed for the four tourism regions of Wales using only data available during the Covid-19 pandemic. The superior performance of the SDM advocates in favour of adding it to the battery of techniques used in tourism forecasting during and after a major economic crisis. Keywords: Demand forecasting, Economic crisis, Regional, State-dependent model

RISE: Recovery-Informed Strategy Enhancement in Post-COVID-19 Chinese Outbound Tourism Forecasting

Presenter: Taozhu Ruan (Central University of Finance and Ecnomics)
Author(s): Taozhu Ruan;Feng Li;Yiming Zhong;Yuchen Xue;Kexin Shi

This paper introduces RISE (Recovery-Informed Strategy Enhancement), a three-stage framework for forecasting the recovery of Chinese outbound tourism following the COVID-19 crisis. In the first stage, we establish a baseline forecast by training 17 time series models on pre-pandemic data and combining them using multiple forecasting strategies. The second stage extends the existing outbound tourism data by estimating a reference series based on external indicators, such as search index and flight information. In the final stage, three recovery curve assumptions—linear, quadratic, and logistic—are applied to adjust the baseline forecasts. By integrating up-to-date information and employing combination techniques throughout all stages of forecasting, RISE improves the accuracy and robustness of tourism recovery predictions. Our framework achieved first place in the point forecasting track and third place in the interval forecasting track of the Tourism Forecasting Competition II, demonstrating its effectiveness in real-world forecasting scenarios.

Recovery of the Chinese Outbound Travel Market from the Pandemic

Presenter: Hanyuan Zhang (The Hong Kong Polytechnic University)
Author(s): Hanyuan Zhang;Anyu Liu;Vera Shanshan Lin;Curley Ying Liu;Xinyang Liu

The unprecedented outbreak of COVID-19 profoundly disrupted international tourism, triggering a dramatic structural break in demand patterns. As many markets gradually enter the post-pandemic recovery phase, accurately forecasting the rebound of outbound travel remains critical for tourism stakeholders and policymakers. This study proposes a combined judgmental forecasting approach to predict Chinese outbound travel recovery for the period from August 2023 to July 2024, focusing on 20 key destinations. In the initial stage, baseline forecasts were generated using ten individual models, including time series, econometric, and artificial intelligence approaches, augmented by event and trend dummies to address pandemic-induced irregularities. This was followed by the judgmental selection of the most accurate model outputs, which were merged via simple averaging. Finally, to accommodate real-world considerations such as flight capacities and policy interventions, additional judgmental adjustments were introduced. This multi-stage process leverages both statistical and contextual insights to enhance forecasting accuracy. The results underscore that destinations in closer proximity to China, including Macao, Hong Kong, Thailand, and Korea, are predicted to recover faster due to shorter travel distances and fewer policy barriers. In contrast, destinations such as Canada, the United States, and Australia may experience delayed recoveries, attributed to less favorable policies and higher travel costs. The obtained forecasting outcomes hold significant relevance for the tourism industry, destination management, and marketing organizations in China, facilitating strategic and operational planning in the

post-pandemic era.

Value-added Analysis of Electricity Price Forecasts

Presenter: Tao Hong (University of North Carolina at Charlotte)
Author(s): Tao Hong

In the literature of electricity price forecasting, only a few papers have been found on the topic of the economic values of the forecasts. We propose an optimization framework to assess the economic value of electricity price forecasts via battery-based energy arbitrage. The importance of battery energy storage systems (BESS) will continue to increase with the integration of renewable energy sources and therefore provides an excellent application of electricity price forecast. In this presentation, we will demonstrate that electricity price forecasts have a noticeable value in battery-based energy arbitrage, while the (statistical) error measures of electricity price forecasts are not necessarily reflective of their value.

Statistical and Economic Evaluation of Forecasts in Electricity Markets: beyond RMSE and MAE

Presenter: Katarzyna Maciejowska (Wroclaw University of Science and Technology)
Author(s): Katarzyna Maciejowska;Bartosz Uniejewski;Arkadiusz Lipiecki

In recent years, one could observe a rapid development of forecasting methods leading to an increase in the accuracy of predictions. In the literature, point forecasts are typically assesed with measures such as Root Mean Squared Errors or Mean Absolute Errors. Although suitable for statistical evaluation, they do not represent well the economic value. In this research, we consider a decision problem of a storage unit. It needs to decide when to charge and discharge the battery using day-ahead forecasts of electricity prices. Conditional on the choice, the trade brings different revenues and hence impacts the financial outcome of the utility. In order to examine the relationship between the accuracy of electricity price forecasts and economic outcomes, we use 24 models/model specifications calibrated with data coming from windows of different sizes. Hence, for each day we use 200 different predictions to support the decision of the battery manager. The forecasts are first evaluated with six statistical scores that go beyond RMSE and MAE and represent different properties of predictions and prediction errors. Finally, the correlation between statistical measures and profit is calculated. The results indicate that RMSE and MAE are poorly related to the revenue level. On the contrary, scores that evaluate the daily distribution of prices are much better predictors of profits and hence should be used to select an optimal model or included in the estimation process.

Almost Exact Matching for Interpretable Day-ahead Electricity Price forecasting

Presenter: Fuyang Jiang (KU Leuven)
Author(s): Fuyang Jiang;Hussain Kazmi

Reliable access to electricity underpins much of modern civilization. In recent years, many different electricity markets (ranging from long-term forwards and futures to short-term day-ahead and intra-day markets) have been established in countries across the world to ensure security of supply and provide users with price signals that reflect moments of scarcity and abundance. The day-ahead electricity market is the most actively traded market in Europe and provides price signals on an hourly basis. To make the most of price arbitrage on this market, traders and retailers typically need to forecast prices before scheduling their own (flexible) demand and generation. Recent years have seen an increasing shift towards more data-driven techniques to forecast electricity market prices. This is driven by greater availability of data and compute resources, and often leverages machine learning algorithms, including regularized linear models, tree-based methods, and neural network variants. Although these models often achieve state of the art performance in terms of predictive accuracy, their industrial application is constrained by their black-box, non-interpretable nature. Unfortunately, existing (post-hoc) model explanation tools such as permutation and feature importance, and the increasingly popular SHAP and LIME libraries provide only an incomplete view of the model's internal functioning, and suffer from their own shortcomings (e.g. tunable hyperparameters, issues with assumptions around feature independence and multicollinearity etc.). In this paper, we propose an interpretable-by-design

matching-based method to forecast electricity prices and apply it to the case of day-ahead electricity markets. Known issues of matching (high computational expense and lack of high-quality matches) are offset by prioritizing matches on important features based on domain expertise and existing models. We also explore model sensitivity to different parametrizations, including the choice of input features, the calibration window length, and the normalization schemes. Preliminary results show that these methods can offer comparable forecast accuracy while remaining interpretable, thus effectively overcoming the issues introduced by post-hoc explanations of black-box methods. The sensitivity analysis also yields new insights into why ensembling different models improves price forecast accuracy.

M-SALRL: A Novel Enhanced Residual Learning Approach For Crude Oil Prices Forecasting With Mixed-Frequency Structure

Presenter: Jian Wu (Chongqing Normal University)
Author(s): Jian Wu;Xianning Wang;Yuke Chen

Violent fluctuations in crude oil prices are intrinsically linked to global economic growth, energy security, and policy formulation, thereby necessitating the anticipation of their trends in advance. These price changes are inherently marked by high nonlinearity and complex evolution patterns. Whereas existing prediction models exhibit pronounced limitations in addressing nonlinear dynamics, capturing temporal regularities, and integrating mixed-frequency data, this study introduces an innovative mixed-frequency enhanced residual learning prediction framework—M-SALRL (MIDAS-SVR-Augmented LSTM Residual Learning). The model is composed of three pivotal modules: 1) constructing a daily-monthly frequency domain information fusion architecture to effectively integrate 10 types of high- and low-frequency indicators, such as crude oil prices and the energy uncertainty index; 2) designing an augmented layer to capture the temporal regularities of nonlinear fluctuations in the crude oil price sequence; and 3) incorporating a residual learning mechanism to decompose the long-term trend components within the crude oil price sequence, coupled with an enhanced stepwise Bayesian optimization approach (I-BO) for parameter optimization. Empirical analyses conducted on a real dataset spanning 156 months (4836 daily observations) of Brent crude oil futures reveal that the mean squared error (MSE) of M-SALRL is diminished by an average of 67.3% relative to fifteen benchmark models. Notably, the prediction error improvement during periods of extreme price volatility amounts to 62.5%, markedly surpassing extant methodologies. Furthermore, during periods of extreme uncertainty, such as the COVID-19 pandemic, the model demonstrates robust adaptability across diverse data frequencies and volatility regimes. This study not only provides a fresh perspective on crude oil price prediction but also furnishes a dynamic quantitative tool of significant value for investment decision-making and energy security early warning.
Keywords: Crude oil price prediction; Mixed-frequency data sampling (MIDAS); Long short-term memory neural network (LSTM); Support vector regression (SVR); Residual learning mechanism (Residual Learning)

Accounting for the Risk of Model Misspecification in Model Selection

Presenter: Nikolaos Kourentzes (University of Skövde)
Author(s): Nikolaos Kourentzes;Ivan Svetunkov

Forecasting models are expected to provide an objective projection of a time series into the future. The forecast uncertainty, as encapsulated by the predictive distribution, should be accounted for by decision makers in their decisions. When they use the mean forecast, or a specific quantile, they make a statement about their risk preferences. Forecast uncertainty is the composite of the uncertainty due to innovations in the time series and any uncertainty due to model misspecification. Since one risk is extrinsic while the other one is intrinsic under the merit of the forecasting team, equal risk preferences from decision makers are unlikely. However, in practice, for unknown data generating processes it is challenging to separate these various sources of uncertainty. Instead, we focus on model selection and propose to explicitly account for the model misspecification risk when selecting from a pool of models. We develop a procedure to achieve this and demonstrate its effect on the characteristics of forecasts.

Identifying Preference from Response Times and a Reference Choice

Presenter: O-Chia Chuang (Hubei University of Economics)
Author(s): O-Chia Chuang;Xiaojun Song;Yinyue Zhang

The present study employs econometric theories and empirical applications to investigate the relationship between decision-makers' response times and their underlying preferences. Our focus is on binary choices with a single reference option, with the objective of inferring the decision-makers' preferences from their response times. To this end, we construct tests targeting the inequality condition proposed by Al´os-Ferrer, Fehr, and Netzer (2021), and the tests are pivotal to the unknown dependence structure of estimation effects. Our empirical study demonstrates that the adjustment of the significance level has a direct impact on the identification of preference relationships. When α is set to 0.05, only 10.37% of the total binary choices can be predicted, with a high prediction accuracy of 80.75%. By increasing α to 0.5, all binary choices can be predicted, but the accuracy drops significantly to 64.42%. Consequently, the optimal significance level is contingent upon the specific circumstances. Our tests demonstrate that response time data can be employed more extensively and rigorously to infer decision-makers' preferences.

Scaling AI Demand Forecasting solution to Enterprise-graded: Reduce the Operational Costs with proper technical designs - A lesson from Unilever.

Presenter: Duc Lai (UNILEVER VIETNAM INTERNATIONAL COMPANY LIMITED)
Author(s): Duc Lai;Tilak Singh;Anu Thomas;Phu Le

This paper addresses the critical gap between theoretical accuracy in demand forecasting algorithms and their practical implementation costs in corporate supply chain management. While statistical precision remains the primary focus of academic research, our investigation reveals significant hidden expenses stemming from computational requirements, processing delays, and vendor dependencies that can negate the benefits of improved forecasting accuracy. We present a novel comprehensive framework for evaluating the total cost of forecasting systems, incorporating computational overhead, maintenance requirements, and vendor-related expenditures. Additionally, we introduce the use of Ray Parallel framework with computing utilization metrics that extend evaluation criteria beyond accuracy, providing a more balanced assessment of overall system efficiency. Our findings yield actionable technical recommendations for optimizing forecasting infrastructure to achieve cost-efficiency without compromising reliability. This research equips organizations with practical tools to make informed decisions that balance accuracy improvements against operational costs, ultimately enabling significant savings across demand forecasting processes while maintaining supply chain performance.

Forecasting Automobile Demand: Spatiotemporal and Hierarchical Modeling, Life Cycle Dynamics, and User-generated Online Information

Presenter: Tom Nahrendorf (Technical University of Munich, Logistics and Supply Chain Management)
Author(s): Tom Nahrendorf;Stefan Minner

This study forecasts monthly automobile demand in a multi-product, multi-market environment. We focus on critical factors in the automotive industry, including declining demand over a product's life cycle, spatiotemporal correlations across regions and time periods, and the complex multi-level structure of demand data in product and market hierarchies. To conduct our analysis, we collected historical demand data from a leading German car manufacturer in the premium automotive sector. We also utilized user-generated online information from an online configurator, supplemented with relevant public macroeconomic indicators. We employ machine learning-based multi-step time series forecasting methods, including a hybrid of exponential smoothing and recurrent neural networks, as well as LightGBM ensembles with various pooling strategies. Additionally, we provide stochastic forecasting extensions using quantile regression and hierarchical reconciliation approaches, such as bottom-up, top-down, and optimal methods. These methods are benchmarked against traditional statistical models, including naïve approaches, moving averages, exponential smoothing, Holt-Winters, ARIMA, and Prophet. We observe a 25.53% improvement in forecast accuracy compared to the best traditional benchmark and an 8.55% enhancement relative to the internal company benchmark, measured by the weighted root mean squared scaled error. Our analysis reveals life cycle curves characterized by exponential growth and decline phases, with multiple peaks during the maturity phase, challenging classical product life cycle models. Therefore, we identify weighted life cycle features as a primary demand indicator, followed by product stock information and configurator website visits, which serve as early indicators of consumer preferences.

Demand Forecasting for New Items

Presenter: Ruben Crevits (OMP)
Author(s): Ruben Crevits

Forecasting the demand of a new item is a challenging and time consuming task for demand planners. How to forecast an item without historic sales data? Classic time series forecasting methods don't address this problem. In this paper we propose a method to still make a forecast, solely based on non-sales features like product properties, customer properties, a promotion plan, calendar events. We illustrate the intuition behind the method with a simulation exercise. We evaluate accuracy and bias on a real data set. We show cases of a reasonable forecast, where there is room for improvement, and where it is not possible to make a good forecast given the data.

RASOR: A Retrieval-Augmented Semiotic Recursion Framework for Adaptive Time-Series Forecasting

Presenter: Ryan Fattini (Farabi Innovations)
Author(s): Ryan Fattini;Ryan Young

Retrieval-augmented generation (RAG) has enhanced large language models by providing verifiable contextual information that reduces hallucinations and improves inference. We extend this concept to time-series forecasting with RASOR, a Retrieval-Augmented Semiotic Recursion framework that integrates natural language processing with statistical modeling, synthetic data generation, and event-driven causal components. At the framework's foundation is a counterfactual baseline, defined by a "business-as-usual" scenario against which real-world events (e.g., promotions or weather anomalies) are modeled. These events are captured through linear or generalized linear components, whose parameters (e.g., via ordinary least squares or exponential-family link functions) encode localized behaviors. Each parameterized function is stored in a semantically labeled repository, enabling embedding-based retrieval for new scenarios, preserving interpretability and limiting parameter inflation. We introduce a semiotic kernel to align probabilistic forecasts with user intent, weighting retrieved components by semantic similarity. This process recurses over meaning-bearing functions, producing refined probability distributions that adapt to emergent conditions without exhaustive retraining. An ensemble mechanism further incorporates user feedback by weighting scenario-specific forecasts according to historical accuracy and bias. Deployed since January 2024 in a retail setting, RASOR demonstrates consistent reductions in mean squared error (MSE) and symmetric mean absolute percentage error (sMAPE) compared to standard statistical baselines and manual overrides over eight months of operations. RASOR provides a flexible and interpretable framework that unifies context-aware retrieval and formal statistical grounding to advance adaptive time-series forecasting.

Optimal Starting Point for Time Series Forecasting

Presenter: Yiming Zhong (Central University of Finance and Economics)
Author(s): Yiming Zhong;Yinuo Ren;Guangyao Cao;Feng Li;Haobo Qi

Recent advances on time series forecasting mainly focus on improving the forecasting models themselves. However, when the time series data suffer from potential structural breaks or concept drifts, the forecasting performance might be significantly reduced. In this paper, we introduce a novel approach called Optimal Starting Point Time Series Forecast (OSP-TSP) for optimal forecasting, which can be combined with existing time series forecasting models. By adjusting the sequence length via leveraging the XGBoost and LightGBM models, the proposed approach can determine the optimal starting point (OSP) of the time series and then enhance the prediction performances of the base forecasting models. To illustrate the effectiveness of the proposed approach, comprehensive empirical analysis have been conducted on the M4 dataset and other real world datasets. Empirical results indicate that predictions based on the OSP-TSP approach consistently outperform those using the complete time series dataset. Moreover, comparison results reveals that combining our approach with existing forecasting models can achieve better prediction accuracy, which also reflect the advantages of the proposed approach.

Do Global Forecasting Models Require Frequent Retraining?

Presenter: Marco Zanotti (University of Milano Bicocca)
Author(s): Marco Zanotti;Matteo Pelagatti

In an era of increasing computational capabilities and growing environmental consciousness, organizations face a critical challenge in balancing the accuracy of their forecasting models with computational efficiency and sustainability. Global forecasting models, which leverage data across multiple time series to improve prediction accuracy, lowering the computational time, have gained significant attention over the years. However, the common practice of retraining these models with new observations raises important questions about the costs of producing forecasts. Using ten different machine learning and deep learning models, we analyzed various retraining scenarios, ranging from continuous updates to no retraining at all, across two large retail datasets. We showed that less frequent retraining strategies can maintain the forecast accuracy while reducing the computational costs, providing a more sustainable approach to large-scale forecasting. We also found that machine learning models are a marginally better choice to reduce the costs of forecasting when coupled with less frequent model retraining strategies as the frequency of the data increases. Our findings challenge the conventional belief that frequent retraining is essential for maintaining forecasting accuracy. Instead, periodic retraining offers a good balance between predictive performance and efficiency, both in the case of point and probabilistic forecasting. These insights provide actionable guidelines for organizations seeking to optimize forecasting pipelines while reducing costs and energy consumption.

Exploring the Advantages and Applications of Large Language Models in Long-Term Stock Price Trend Prediction

Presenter: Thithuy Bui (University of Chinese Academy of Sciences (UCAS))
Author(s): Thithuy Bui;Zhengjun Zhang

With the rapid advancement of artificial intelligence and big data technologies, along with the ongoing evolution of the global economic landscape, the accuracy of stock price prediction has become increasingly critical. In the financial sector, the application of large language models (LLMs) is gradually expanding. This study innovatively integrates stock market data with large language models by constructing the Time-LLM framework, and effectively enhances stock price prediction accuracy through the use of prompt engineering. Taking the top five companies by market capitalization in both the United States and China, as well as a major Chinese market index as research subjects, the study compares the predictive performance of four traditional statistical models, four machine learning models, and the large language model GPT-4o-mini. The findings reveal that the GPT-4o-mini-based LLM significantly outperforms traditional statistical and machine learning models in stock price forecasting for four out of the top five U.S. companies by market capitalization, as well as all three components of the Chinese market index, while traditional statistical models performed best for one of the five U.S. companies.In the stock price prediction of the top five companies by market capitalization in China, the large language model performed best for two companies and second-best for another two, while machine learning models performed best for three companies. We find that the large language model GPT-4o-mini is capable of delivering generalized predictions tailored to the specific characteristics of individual companies' stock prices. This research not only highlights the high accuracy and low error rate of GPT-4o-mini in stock price prediction, but also offers potential technical pathways and theoretical foundations for the application of large language models in the financial domain.

Research on the Effectiveness of ESG Screening Strategies in China's A-share Market—From the Perspective of Industry Heterogeneity

Presenter: Zheng Sinian (University College Dublin)
Author(s): Sinian Zheng;Valerio Poti;Alessia Paccagnini

Despite the growing adoption of Environmental, Social, and Governance (ESG) strategies in global financial markets, the variability in their effectiveness across different industrial sectors and stages of corporate development remains inadequately explored. Particularly in China's A-share market, there is a significant gap in understanding how industry-specific characteristics influence the outcomes of ESG screening over long periods. This study addresses

this gap by examining the impact of ESG screening strategies on the financial performance of companies within various sectors of China's A-share market from 2012 to 2022. Using Bloomberg ESG scores to construct long-short portfolios, this research employs a robust methodology including the Discounted Cash Flow (DCF) model to dissect the differential effects of ESG strategies on cash flows, valuation, and risk management across growth, maturity, and decline life cycles. The findings reveal that ESG screening consistently generates positive abnormal returns, with negative screening strategies notably outperforming positive ones. However, the impact varies considerably across industries: ESG strategies are less effective in growth-stage companies compared to those in mature or declining phases. This heterogeneity is primarily due to how different industry characteristics modulate the influence of ESG factors on operational cash flows, market valuations, and risk levels. These results not only confirm the critical role of tailoring ESG integration to specific industry contexts but also highlight the necessity for investors and policymakers to adapt ESG strategies based on the distinctive attributes of each sector. Furthermore, the study underscores the urgent need for ongoing research to unravel the complex interactions between ESG practices and financial performance, particularly in emerging markets like China. This research advances our understanding of sustainable investment and aids in the refinement of ESG evaluation frameworks, contributing significantly to the broader field of sustainable finance. Keywords: Asset Pricing Models, ESG Screening Strategies, China's A-share Market, Industry Heterogeneity, Risk Management, Investment Strategies, Emerging Markets

The Impact of Environmental, Social and Governance (ESG) Scores on Financial Performance: Evidence from Malaysia

Presenter: Chee Pung Ng (INTI International University)
Author(s): Chee Pung Ng

With the ongoing environmental crisis and the growing emphasis on Corporate Social Responsibility (CSR), Environmental, Social, and Governance (ESG) disclosure has gained significant traction. This study investigates the impact of ESG scores on financial performance across three sectors in Malaysia from 2013 to 2021. ESG performance is evaluated using the Refinitiv Eikon ESG disclosure score, while financial performance is measured by Return on Assets (ROA). These panel data were analysed with pooled ordinary lest square, fixed effect model, random effect model. Diagnostic checking tests such as normality test, heteroscedasticity test, and multicollinearity test were conducted in this study. Panel regression analysis demonstrates a significant relationship between ESG and financial outcomes. The findings offer strategic insights for firms, investors, and policymakers, encouraging greater transparency in ESG reporting and reinforcing the financial benefits of sustainable corporate practices.

Undercoverage of Conformal, Normal Prediction Intervals

Presenter: Casey Lichtendahl (Google)
Author(s): Casey Lichtendahl

Backtest errors are often used to construct conformal, normal prediction intervals. The width of these intervals is assumed to be proportional to the estimated standard deviation of the backtest errors. However, this assumption misses the uncertainty in the estimated variance of the backtest errors. When backtest errors are independently and normally distributed, their estimated variance follows a scaled chi-squared distribution. Unfortunately, backtest errors from longer-horizon forecasts can be strongly positively serially correlated, and their estimated variance will not follow a scaled chi-squared distribution. Nonetheless, if we estimate the effective (independent) sample size of our backtest errors, we can use this distribution to approximately correct the undercoverage of conformal, normal prediction intervals. Our correction is based on the assumption that the estimated error variance is distributed as a scaled chi-squared random variable with degrees of freedom equal to the effective sample size of our backtest errors. The correction results in a Student's t forecast distribution with closer to ideal coverage.

Multivariate Neural Network for Joint Distribution Forecasting

Presenter: Pablo Montero Manso (Google)
Author(s): Pablo Montero-Manso

While univariate point and marginal distribution forecasting with neural networks has achieved significant success, industrial applications demand multivariate joint distribution forecasts, a challenging extension. Existing methods, such as diffusion models and variational autoencoders, struggle to maintain performance and are difficult to adapt to benefit from business logic knowledge, particularly when incorporating custom-tailored loss functions. We introduce a neural framework designed to overcome these limitations. This framework (1) minimizes architectural changes on already fine-tuned univariate models by simply augmenting them with a compact set of learned multivariate features, and (2) utilizes a Gibbs factorization of the joint distribution, building upon the familiar and tractable marginal multi-horizon distributions used by leading models use and further adaptable by analysts. Our methodology is highly scalable, which is essential for a precise approximation of extreme regions of the joint distribution. These regions are often critical in risk assessment and decision-making across diverse domains.

Quadratic Activation for Quantile Forecasts

Presenter: Weijie Shen (Google LLC)
Author(s): Weijie Shen

Accurate quantile forecast is crucial for many real-world use cases, and training a neural network with quantile loss is often preferred to probabilistic loss because it is distribution free. However, when training multiple quantiles in the same neural network, quantile crossing may happen where lower quantiles end up larger than the higher quantiles. To prevent that, a positive constraint is often put on the quantile gaps, ensuring the monotonicity. There are also other monotonic constraints including increasing uncertainty through horizons. These constraints are often represented as a positive activation layer such as ReLU or Softplus. However, these common choices lead to convergence problems during training such as degenerating to zeros, producing quantiles that are not trustworthy. In this talk, we explore using different activation functions for this monotonicity task and conclude the superior performance of the quadratic activation function. We rationalize the finding and recommend this simple change to stabilize quantile training and make it more efficient.

Forecasting under structural constraints to best support decision optimizations

Presenter: Lijuan Xu (Google)
Author(s): Lijuan Xu

Forecasting is seldom the end of a decision making process, but often the start of it. In many circumstances, forecasting results are not only leveraged by human decision makers for critical insights, but also used by subsequent optimization engines to maximize a target metric or to stay under known constraints. In such cases, forecasting becomes an interesting “constraint optimization” problem where we aim for the best prediction performance while maintaining a specific structural form. In this talk, we will share one such example and the journey we went through to produce a forecast that is accurate yet also optimization friendly.

Value-oriented Forecast Reconciliation for Renewables in Electricity Markets

Presenter: Honglin Wen (Shanghai Jiao Tong University)
Author(s): Honglin Wen;Pierre Pinson

Forecast reconciliation is considered an effective method for achieving coherence and improving forecast accuracy. However, the value of reconciled forecasts in downstream decision-making tasks has been mostly overlooked. In a multi-agent setup with heterogeneous loss functions, this oversight may lead to unfair outcomes, hence resulting in conflicts during the reconciliation process. To address this, we propose a value-oriented forecast reconciliation approach that focuses on the forecast value for individual agents. Fairness is ensured through the use of a Nash bargaining framework. Specifically, we model this problem as a cooperative bargaining game, where each agent aims to optimize their own gain while contributing to the overall reconciliation process. We then present a primal-dual algorithm for parameter estimation based on empirical risk minimization. From an application perspective, we consider an aggregated wind energy trading problem, where profits are distributed using a weighted allocation rule.

We demonstrate the effectiveness of our approach through several numerical experiments, showing that it consistently results in increased profits for all agents involved.

Coherent Density Forecasts for Tourism Demand with Automated Immutability Constraints

Presenter: Quan Wen (Beihang University)
Author(s): Quan Wen;Yanrong Zeng;Fotios Petropoulos;Yanfei Kang

This paper addresses key challenges in tourism demand forecasting by introducing novel methodologies to enhance both point and density (probabilistic) forecasts within a hierarchical time series framework. Tourism demand forecasts, often segmented by geographic locations, travel purposes, and administrative divisions, must satisfy linear aggregation constraints through reconciliation processes. While recent studies have shown that immutability constraints—preserving the forecasts of specific key nodes during reconciliation—can improve the accuracy of point forecasts, point forecasts alone are often insufficient in the tourism sector, where demand is subject to significant uncertainty. In such cases, density forecasts provide crucial insights into forecast uncertainty, enabling decision-makers to manage risks better, allocate resources efficiently, and make more informed strategic decisions. To address these issues, this paper presents two key contributions. First, it proposes two automated methods for identifying immutable nodes during reconciliation: a clustering-based approach that manages the number of immutable nodes at each hierarchical level using a two-stage clustering process, and an optimization-based method that selects immutable nodes by solving a penalized optimization problem. Second, the paper introduces immutability constraints into density forecasting through a novel constrained score optimization model, which determines the density forecast reconciliation weights by optimizing an energy score-based objective function with immutability constraints. Furthermore, the study demonstrates that the constrained optimization problem can be effectively transformed into an unconstrained one. We solve the optimization problems in this paper using the Gurobi optimizer. The effectiveness of this dual approach is demonstrated through a series of Monte Carlo simulations and an empirical analysis of Australian domestic tourism data. The results demonstrate substantial improvements in both flexibility and accuracy across hierarchical levels, providing actionable insights for tourism stakeholders. By addressing the inherent uncertainties and structural complexities of tourism demand forecasting, this methodology offers a robust solution for optimizing resource allocation, managing risk, and informing strategic decision-making in tourism management.

Immutable Forecasts for Hierarchical Time Series: An Approach through Trace Minimisation

Presenter: Shanika Wickramasuriya (Monash University)
Author(s): Shanika Wickramasuriya;George Athanasopoulos;Ahmed Samah

Forecast reconciliation for hierarchical time series involves generating base forecasts for each series in the structure, then adjusting them to ensure coherence across the aggregation structure. In some applications, certain base forecasts must remain unchanged (immutable) during this process. This paper introduces a novel methodology for handling immutable forecasts by framing the problem as a constrained optimisation, minimising the total variance of reconciled forecast errors while enforcing equality constraints. Since forecasts of all series cannot be kept immutable simultaneously, we derive conditions for selecting a valid set of immutable forecasts that maintain coherency after the reconciliation process. An empirical application using Australian domestic tourism data shows that the proposed method outperforms an existing alternative.

Dynamic Characteristics Analysis and Precision Enhancement in CPI Hierarchical Forecasting for China

Presenter: Han Jin (Jilin University)
Author(s): Han Jin;Han Liu

Accurate Consumer Price Index (CPI) forecasting is vital for macroeconomic policy formulation, business decision optimization, and inflation expectation management, playing an irreplaceable role in ensuring people's livelihood, stabilizing society, and aiding national economic accounting.The CPI has a complex structure, covering numerous

goods and services categories. Traditional CPI forecasting methods overlook the differences in consumption patterns across sub-indices, failing to capture the transmission differences between volatile categories like food and energy and the core CPI. This shortcoming can lead to delayed policy responses. The hierarchical forecasting method enhances both prediction accuracy and economic interpretability through deconstructing price formation mechanisms. This study employs a Hierarchical Time Series (HTS) forecasting model with dynamic weighting. Using a two-way coordination mechanism between hierarchy levels, the model enhances forecasting accuracy, dynamically captures the contribution of sub-index CPI to the overall index, and accurately identifies the sources and characteristics of inflationary pressures. The empirical results demonstrate that: (1) The dynamic-weighted hierarchical time series model significantly enhances forecasting accuracy. By implementing a bidirectional reconciliation mechanism to ensure consistency between aggregate and sub-indices, the model reduces prediction errors by 15%, 22%, and 28% for short-term, medium-term, and long-term horizons respectively, effectively mitigating systemic biases inherent in conventional forecasting approaches. (2) The CPI dynamic weight heatmap analysis reveals pronounced cyclical variations: during inflationary periods, food and energy categories exhibit substantially elevated contribution shares to the aggregate index (reaching 32% and 28% respectively), whereas services dominate during stable phases, accounting for 45% of the total variation. (3) The hierarchical forecasting can precisely identify the main drivers and transmission characteristics of inflation pressure two to three months in advance, boosting policy response efficiency by 30% and offering more precise decision-making support for policymakers. This study optimizes policy response efficiency and coordinates the regulation of people's livelihood commodity prices to assist in the timely adjustment and implementation of macroeconomic policies.

Data-driven Models for Energy Forecasting, without the Data

Presenter: Hussain Kazmi (KU Leuven)
Author(s): Hussain Kazmi

The ongoing energy transition is necessitating better, more sophisticated forecasting methods for stable electricity grid operation. With increasing proliferation of renewables and demand electrification, this need will grow further in importance in the coming years. To address these challenges, the modern forecasting stack increasingly relies on potentially data-intensive and sample-inefficient machine learning techniques such as neural networks. In other words, these models require a lot of data and compute power to learn accurate mappings between input and output variables. This observational data is not always available, either because of commissioning of new assets in new locations (e.g. offshore wind farms in the North Sea, solar PV systems in impoverished communities), or data acquisition costs (including privacy and security concerns etc.). In this talk, I will discuss how relatively young themes in forecasting literature, including transfer and physics-informed learning, can help alleviate many of these concerns. More concretely, starting with the concept of sim2real, I will cover how pre-training models using simulated data from an existing simulator can be beneficial in real-world settings with transfer learning. Following this, I will relax the assumption of the presence of a perfect simulator in two key ways: (1) where only a mis-specified or imperfect simulator of the process being modelled is available, and (2) where no simulator is available at all. The former formulation, tricky as it may sound, lends itself well to domain randomization and adaptation techniques. The latter, on the other hand, is more challenging still but can benefit greatly from collaborative learning where this is a possibility. When this is not possible, other means of biasing model predictions (in a principled manner) can be beneficial as a last resort. Drawing on several energy forecasting case studies for generation and demand, I will show that these innovations can allow for closer integration between different forecasting paradigms, and lead to models that significantly outperform their existing (black- or white-box) counterparts at a fraction of the data expense. This family of algorithms also represent a marked improvement over classical hybrid or grey-box models, which are typically constrained in terms of model parameters, as opposed to the over-parametrized models common in these formulations.

Composite Sentiment Index based Crude Oil Prediction: a Large Language Model approach

Presenter: Yishuai Li (Hunan normal university)
Author(s): YiShuai Li;Kaijian He;YingChao Zou

With the rapid development of technology, participants in the crude oil market are exposed to the increasingly

complex and versatile risky factors, They are represented in different big data such as the texutal data and time series data. Traditional econometric and time series models have difficulty in modeling new forms of big data such as textual data. Sentiment analysis has been proposed as an effective technique to capture the market impacts of different risky factors, represented in the textual form. In this paper, a new crude oil forecasting model with composite sentiment index based on Large Language models has been proposed.The composite sentiment index has been consutrcted as the averages of a series of individuals sentiment scores estimated with different Large Language Model. This composite sentiment score correspond to a market consensus and is incorporated in the armax based crude oil forecasting model. Results from comprehensive empirical evaluation of different models show that composite sentiment index based forecast model produce crude oil price forecasts with a higher level of accuracy. Experiment results suggested that the performance is sensitive to the choice of large language model.

The Universality and Predictability of Technology Adoption

Presenter: Benjamin Wagenvoort (University of Oxford)
Author(s): Benjamin Wagenvoort;Joel Dyer;François Lafond;J. Doyne Farmer

We develop and test a method to make distributional forecasts of technology diffusion based on the Gompertz and Logistic curves. First, we construct the hitherto largest database of technology adoption. We include over 150 technology adoption curves across various industries and centuries. Second, we show that most of the technologies have a r-squared over 0.90 when fitting the 3-parameter symmetric logistic equation or Gompertz curve. Third, based on the autocorrelation properties of the data, we propose stochastic versions of the logistic and Gompetz curves. Finally, we use Bayesian inference to make well-calibrated distributional forecasts. When we backtest our forecasting methodology at 10% diffusion, the mean absolute percentage error is consistently below 0.40 across the remaining 90% of the adoption curve. We apply our method to renewable energy technologies. We find that global Solar PV capacity is likely to surpass the net-zero requirements set by the International Energy Agency (IEA) of 14-16.5 TW by 2050.

Joint First-Order Learning of Smoothing and Model Parameters in Hybrid GAM-FNN Forecasting

Presenter: Monika Zimmermann (Universität Duisburg-Essen)
Author(s): Monika Zimmermann;Florian Ziel

Generalized Additive Models (GAMs) are widely used in time series forecasting due to their interpretabil- ity and computational efficient implementation within the R-package mgcv. While efficient in capturing up to second-order interactions, GAMs become less scalable for modeling higher-order interactions com- pared to Feedforward Neural Networks (FNNs), making hybrid GAM-FNN models appealing. However, existing hybrid models fix smoothing parameters, risking overfitting. To address this, we propose a novel efficient first-order estimation criterion that learns smoothing parameters jointly with model parameters. Our approach, initially discussed in a Gaussian setting, minimizes an estimate of the Kullback–Leibler (KL) divergence using the joint rather than conditional likelihood in the Linear Mixed Model (LMM) view. Numerical experiments show that gradient-based optimization with this criterion yields results comparable to common smoothing parameter estimation methods (REML, GCV, and UBRE). Unlike UBRE and GCV, which have shallow score functions near their minima, our criterion results in more pronounced and stable local minima under resampling. A forecasting study on high-resolution electricity load data demonstrates that a hybrid GAM-FNN model trained with our method significantly outper- forms a REML-estimated GAM in terms of RMSE underscoring its practical advantages.

Spatiotemporal Solar Forecasting Using a Single Sky Imager and Satellite Data

Presenter: Amar Meddahi (TotalEnergies)
Author(s): Amar Meddahi;Yves-Marie Saint-Drenan;Arttu Tuomiranta;Sebastien Guillon;Philippe Blanc

Solar forecasting is essential to reduce uncertainty in solar power generation, thereby facilitating its integration into energy grids. At intra-hour timescales, ground-based sky imagers enhance forecasting accuracy by providing

high-resolution spatiotemporal estimates of solar irradiance variability. Current spatiotemporal forecasting methods typically rely on additional ground-based instruments, such as remote sensing devices, networks of sky imagers, or dedicated cloud motion sensors. However, installing such equipment poses challenges for photovoltaic (PV) plants, especially due to limited network connectivity, energy constraints, and maintenance complexities. In this work, we introduce a spatiotemporal solar forecasting framework that utilizes only a single on-site sky imager and readily accessible satellite weather data. The proposed method estimates cloud base height (CBH) by integrating data from the sky imager with satellite observations, removing the necessity for additional ground-based instrumentation. This reduction in hardware dependency expands the practical applicability of spatiotemporal solar forecasting at PV installations. The forecasting framework leverages the kinematic relationship connecting the effective linear cloud speed, obtained from satellite-derived Atmospheric Motion Vectors (AMV), and its corresponding pixel-based speed derived from the ground-based sky imager, to estimate cloud height. This kinematic relationship allows for the effective use of satellite-derived AMV wind speeds as linear cloud speed, while standard cloud-motion tracking algorithms yield pixel cloud speed. The estimated height can then be integrated into existing spatiotemporal forecasting methods, where previously an additional ground-based sensor was necessary. We evaluate the proposed framework through two analyses. First, we assess the uncertainty of our cloud height estimation by comparing it to ground-based measurements and to the ECMWF-IFS numerical model, which represents a baseline alternative when additional on-site sensors are unavailable. Second, we analyze the impact of height estimation errors on the accuracy of solar irradiance forecasts on the ground, noting that the effect depends also on solar position. The evaluation utilizes experimental data from the SIRTA observatory and wind speed fields derived using the EUMETSAT MSG - 0 degree AMV algorithm.

Adaptive Probabilistic Forecasting for Wind Energy Based on Generalised Logit Transformation and Bayesian Method

Presenter: Tao Shen (University of Glasgow)
Author(s): Tao Shen;Jethro Browell;Daniela Castro-Camilo

Wind power plays an increasingly significant role in achieving the 2050 Net Zero Strategy. Despite its rapid growth, its inherent variability presents challenges in forecasting. Accurately forecasting wind power generation is one key demand for the stable and controllable integration of renewable energy into existing grid operations. This paper proposes an adaptive probabilistic forecasting method that combines the generalised logit transformation with a Bayesian approach. The generalised logit transformation processes double-bounded wind power data to an unbounded domain, facilitating the application of Bayesian methods. A novel adaptive mechanism for updating the transformation shape parameter is introduced to leverage Bayesian updates by recovering a small sample of representative data. Four adaptive forecasting methods are investigated, evaluating their advantages and limitations through an extensive case study of data over 100 wind farms ranging four years in the UK. The methods are evaluated using Continuous Ranked Probability Score (CRPS) metrics, supplemented by reliability and sensitivity analysis. Results indicate that the proposed Bayesian method with adaptive shape parameter updating outperforms benchmarks, yielding consistent improvements in CRPS and forecast reliability. The method effectively addresses uncertainty, ensuring robust and accurate probabilistic forecasting which is essential for grid integration and decision-making.

Long-term Forecasting Model for the State of a Zonal Electric Power System.

Presenter: Eugene Talygin (SKM.PRO)
Author(s): Eugene Talygin

Zonal electric power systems (ZEPS) can be schematically represented as a set of interconnected zones with a given set of generating equipment and foreknown consumption. Each zone has a certain set of generating equipment with known parameters, as well as a set of electricity consumers whose total demand is known. The zones are interconnected by lines that provide electricity exchange. This exchange is necessary to maintain a balance between consumption and generation in each zone and in the entire system. The main feature of ZEPS is the fact that the transmission of electricity from a generator to a consumer within a zone itself is associated with significantly lower losses than the transmission of electricity between zones. Each generating unit included in the ZEPS is characterized by several basic parameters, the most important of which are minimum capacity, maximum capacity and cost of

electricity production. In this case, the total consumption, considering losses throughout the system, is always less than the maximum generation. Therefore, it can be hard to properly choose the composition of generating equipment in zonal systems. Firstly, the selected equipment should fully cover the existing demand for electricity; secondly, the cost of the generated electricity should be as low as possible. In the context of this study, a long-term forecasting system for the state of a zonal electric power system based on mixed integer programming was developed. The input to the system under consideration is statistical data on consumption and reference data on the parameters of each generating unit. In addition, the approximate composition of the generating equipment is assumed to be known. Then the problem of optimizing the composition of the generating equipment occurs in 2 stages - at the first, the mixed integer problem of selecting the included equipment is solved, and at the second, the capacity of the selected equipment is optimized. Then the marginal prices of generation are obtained as dual variables to the balance constraints. The developed model allows building long-term forecasts for a period of decades.

A Novel Concept Bottleneck Generation Model for Explainable Financial Risk Prediction

Presenter: Zhao Wang (Hefei University of Technology)
Author(s): Zhao Wang;Jianfei Wang;Cuiqing Jiang

With rising global economic uncertainty and increasing volatility in capital markets, investors are increasingly relying on accurate and explainable financial risk prediction methods to optimize investment decisions. However, existing methods typically function as black boxes, lacking interpretability and thus limiting their practical applicability in high-stakes financial contexts. Concepts, as fundamental units of human cognition, encapsulate the essential attributes and intrinsic features of financial phenomena. Concept-based method provides a promising avenue for improving both accuracy and explainability in financial risk prediction tasks. Yet, establishing a precise and transparent mapping from financial texts to concepts and subsequently to risk outcomes remains an open challenge. To address this gap, we propose LLM4CBM, a novel concept bottleneck generation model for explainable financial risk prediction. Specifically, we design a grounded theory-driven concept discovery module, which guides large language models (LLMs) to discover domain-relevant concepts from existing financial texts. By leveraging grounded theory's coding principles—including open coding, axial coding, and selective coding—the module constructs structured concept profiles that capture the essential attributes and relationships of financial phenomena. Furthermore, we design a dual-criterion retrieval-augmented concept annotation module, which empowers LLMs to precisely annotate relevant concepts within new financial texts. By incorporating consistency (between the new text and annotated samples, and between concept profiles and annotated samples) and diversity (among annotated samples) as dual retrieval criteria, the module dynamically selects relevant concept cards and annotated examples to construct effective in-context prompts. Finally, we leverage a logistic regression model that maps identified concepts to financial risk results, ensuring transparent decision pathways. Extensive experiments conducted on stock movement prediction task demonstrate that the proposed LLM4CBM model significantly outperforms baseline methods in both predictive performance and interpretability. This research not only establish a new paradigm for trustworthy AI in financial analytics, but also provides a novel framework for aligning LLMs with human-understandable conceptual reasoning.

Taking the Hints from Auditors: Detecting the Hidden Corporate Frauds Using Expert Knowledge-based Label Noise Filtering Algorithm

Presenter: Xiao Yao (Central University of Finance and Economics)
Author(s): Xiao Yao

Corporate fraud detection faces significant challenges due to label noise caused by delayed fraud disclosure, where hidden fraudulent firms are mislabeled as non-fraudulent. Traditional models relying on financial ratios are vulnerable to manipulation, while existing label noise filtering methods neglect domain-specific knowledge. Leveraging auditors' professional judgment embedded in Key Audit Matters (KAMs) of audit reports, this study proposes a novel active learning-based label noise filtering algorithm that integrates multi-dimensional textual features from KAMs to iteratively refine fraud predictions. Using a sample of China's A-share listed firms, we benchmark the algorithm against a collection of machine learning models and deep learning techniques. Results demonstrate that our approach significantly improves fraud detection accuracy, particularly in identifying undisclosed frauds and reducing misclassification costs. Furthermore, it achieves timely detection of disclosed frauds within the same fiscal

year, offering practical utility for investors and regulators. Our findings advocate for the strategic use of auditors' expert insights in enhancing fraud detection methods, with implications for highlighting the information value of KAM disclosure.

Leveraging Sequential In-Loan Behaviors in Credit Scoring: A Multi-Teacher Distillation Framework with Progressive Knowledge Fusion

Presenter: Zhou Fanyin (Southwestern University of Finance and Economics)
Author(s): Fanyin Zhou;Ke Li;Wei Zhao

Traditional application scoring models in credit risk assessment primarily rely on pre-loan features, overlooking the critical value of sequential in-loan behavioral data (e.g., installment repayment records) that dynamically reflect borrowers' creditworthiness. To address this limitation, we propose a multi-teacher distillation framework that systematically integrates sequential in-loan behaviors into credit risk prediction through three key innovations: (1) a sequential distillation architecture with teacher models targeting stage-specific repayment patterns. (2) a progressive knowledge fusion mechanism using a "proximity-aware" temporal weighting strategy. (3) a parameter-inherited stacking model that enhances robustness by allowing student models to inherit optimized parameters and combining them through ensemble learning. Experiments on real-world credit data demonstrate that our framework significantly outperforms conventional application scoring models. Ablation studies further reveal the superiority of the multi-teacher design over single-teacher approaches, with progressive fusion and parameter inheritance effectively mitigating long-term information loss. This work establishes a practical paradigm for enhancing credit risk models through sequential in-loan behavior analysis, enabling financial institutions to leverage temporal repayment insights without operational disruption.

Uncovering Financial Distress with Textual Risk Disclosures in Annual Reports: Insights from Large Language Models

Presenter: Xiaoqian Zhu (University of Chinese Academy of Sciences)
Author(s): Xiaoqian Zhu;Jianping Li;Hao Sun;Hanlin Jin

Risk factor disclosures in annual reports, which discuss the potential uncertainties faced by companies in textual form, are rarely utilized in financial distress prediction. This paper innovatively employs the large language model (LLM) to extract systematic textual predictors for financial distress from massive textual risk factor disclosures. Based on risk factor disclosures from 22,256 financial reports of 6,598 U.S. publicly listed companies between 2014 and 2023, a total of 24 LLM-based textual predictors are extracted, including two new predictors not included in prior studies. Compared to commonly used textual features, LLM-based textual predictors improve financial distress prediction performance by achieving a significantly higher AUC (Area Under the Curve). Furthermore, following instructions in specially designed prompts, the LLM generates detailed explanations on how to utilize extracted textual predictors to distinguish distressed companies from non-distressed ones, making LLM-based textual predictors highly interpretable. This study highlights the critical role of risk factor disclosures in assessing financial distress risks and demonstrates the powerful textual information extraction capabilities of LLMs.

Communication-Efficient L0 Penalized Least Square

Presenter: Chenqi Gong (Chongqing University)
Author(s): Chenqi Gong;Hu Yang

In this paper, we propose a communication-efficient penalized regression algorithm for high-dimensional sparse linear regression models with massive data. This approach incorporates an optimized distributed system communication algorithm, named CESDAR algorithm, based on the Enhanced Support Detection and Root finding (SDAR) algorithm. The CESDAR algorithm leverages data distributed across multiple machines to compute and update the active set and introduces the communication-efficient surrogate likelihood framework to approximate the optimal solution for the full sample on the active set, resulting in the avoidance of raw data transmission, which enhances privacy and data security, while significantly improving algorithm execution speed and substantially reducing communication costs.

Notably, this approach achieves the same statistical accuracy as the global estimator. Furthermore, this paper explores an extended version of CESDAR and an adaptive version of CESDAR to enhance algorithmic speed and optimize parameter selection, respectively. Simulations and real data benchmarks experiments demonstrate the efficiency and accuracy of the CESDAR algorithm.

Context is Key: A Benchmark for Forecasting with Essential Textual Information

Presenter: Arjun Ashok (University of Montreal)
Author(s): Arjun Ashok

Forecasting is a critical task in decision-making across numerous domains. While historical numerical data provide a start, they fail to convey the complete context for reliable and accurate predictions. Human forecasters frequently rely on additional information, such as background knowledge and constraints, which can efficiently be communicated through natural language. However, in spite of recent progress with LLM-based forecasters, their ability to effectively integrate this textual information remains an open question. To address this, we introduce "Context is Key" (CiK), a time-series forecasting benchmark of 71 tasks spanning 7 domains, that pairs numerical data with diverse types of carefully crafted textual context, requiring models to integrate both modalities; crucially, every task in CiK requires understanding textual context to be solved successfully. We further introduce the Region of Interest CRPS metric (RCRPS) to evaluate context-aided forecasting performance, which prioritizes context-sensitive windows in the prediction and accounts for constraint satisfaction. We evaluate a range of approaches, including traditional statistical models, time series foundation models, and LLM-based forecasters, and propose a simple yet effective LLM prompting method that outperforms all other tested methods on our benchmark. Our experiments highlight the importance of incorporating contextual information, demonstrate surprising performance when using LLM-based forecasting models. Our analysis explores key factors such as the impact of context conditioning and prompting techniques, discusses key failure modes of models, and highlights the relative advantages of traditional and LLM-based forecasters in various settings. This benchmark aims to advance multimodal forecasting by promoting models that are both accurate and accessible to decision-makers with varied technical expertise.

A Novel Self-scaled Approximate l0 Regularization Robust Model for Outlier Detection

Presenter: Song Pengyang (AMSS Center for Forecasting Science, Chinese Academy of Sciences, Beijing, 100190, China)
Author(s): Pengyang Song

Robust regression models in the presence of outliers have significant practical relevance in areas such as signal processing, financial econometrics, and energy management. Many existing robust regression methods, either grounded in statistical theory or sparse signal recovery, typically rely on the explicit or implicit assumption of outlier sparsity to filter anomalies and recover the underlying signal or data. However, these methods often suffer from limited robustness or high computational complexity, rendering them inefficient for large-scale problems. In this work, we propose a novel robust regression model based on a Self-scaled Approximate ℓ_0 Regularization Model (SARM) scheme. By introducing a self-scaling mechanism into the regularization term, the proposed model mitigates the negative impact of uneven or excessively large outlier magnitudes on robustness. The resulting optimization problem is nonconvex; to solve it efficiently, we develop an alternating minimization algorithm grounded in proximal operators and block coordinate descent. We rigorously prove the convergence of the whole sequence, and under the Restricted Isometry Property (RIP) framework, we derive a theoretical upper bound on the estimation error associated with SARM. Empirical comparisons with several state-of-the-art robust regression methods demonstrate that SARM not only achieves superior robustness but also significantly improves computational efficiency. Motivated by both the theoretical error bound and empirical observations, we further design a Two-Stage SARM (TSSARM) framework, which better utilizes sample information when the singular values of the design matrix are widely spread, thereby enhancing robustness under challenging conditions. We provide a detailed applicability analysis of both models and propose practical guidelines for model selection and usage. Finally, we validate our approach on a real-world power load forecasting dataset provided by ISONE. The experimental results show that our model substantially enhances the robustness of load forecasting against adversarial data attacks, which is increasingly critical in the era of heightened data security concerns.

Optimal Parameter-Transfer Learning for Interval-valued Data

Presenter: Haowen Bao (Academy of Mathematics and Systems Science, Chinese Academy of Sciences)
Author(s): Haowen Bao;Yuying Sun

This paper proposes a parameter-transfer learning method for interval-valued data through an optimal model averaging procedure. We develop an innovative approach where a target model incorporates information from auxiliary source models through shared parametric components, while accommodating distinct structural patterns in interval-valued observations across different model specifications. For optimal weight selection, we propose a novel multi-fold forward validation criterion designed for interval-valued data. Under regularity conditions, we establish that the proposed weight estimates achieve asymptotic optimality under model misspecification, and automatically exclude all incorrect source models when the target model is correctly specified. Monte Carlo simulations demonstrate the favorable finite sample properties of our method. An empirical application to asset price forecasting reveals the superior performance of our approach relative to existing methods. Our findings highlight the potential of machine learning techniques in interval-valued data analysis, offering new insights for financial forecasting under uncertainty.

An Interval-based Transformer Approach for Carbon Future Prices Forecasting

Presenter: Chuanmiao Yan (Academy of Mathematics and Systems Science, Chinese Academy of Sciences)
Author(s): Chuanmiao Yan;Yutong Sun;Yao Yue;Yuying Sun;Shouyang Wang

Abstract—We propose a novel interval-based transformer approach called DKformer to forecast quarterly carbon futures prices by treating interval-valued data as inseparable random sets. Our DKformer model leverages the attention mechanism to capture long-term dependencies and complex dynamics in the carbon market. The model integrates the European Economic Policy Uncertainty (EEPU) index and the Climate Policy Uncertainty (CPU) index, utilizing only lags of price values and EPU indices as explanatory variables to effectively capture the impact of policy shifts on carbon futures prices. Empirical tests demonstrate the strong performance of the model in predicting market turning points and responding to various policy change scenarios, including the implementation of the EU ETS Phase 4, the RePowerEU Plan, and the Carbon Border Adjustment Mechanism (CBAM). Furthermore, robustness checks conducted during different forecast periods highlight the merits of our model, further validating the effectiveness in dynamic policy environments. Index Terms—Interval quarterly forecasting, Carbon future prices, Transformer, Policy uncertainty

Forecasting Lowest VaR of Bitcoin: An Interval-valued Data Model with Multiple Influencing Pridictors

Presenter: Dingxuan Zhang (Chang'an University)
Author(s): Dingxuan Zhang

This study proposes a novel interval-valued modeling framework for forecasting Value at Risk (VaR) associated with Bitcoin's lowest price (hereafter termed lowest VaR). We integrate the Autoregressive Conditional Interval (ACI) model with a bootstrap technique and decay factor to develop confidence interval estimators: the Bootstrap Eclipse (BE) and Low Quantile (LQ) methods. Empirical results demonstrate the superior performance of our interval-based approach compared to conventional point-valued methodologies. The analysis also reveals distinct application domains for BE and LQ methods under varying risk scenarios. The parameter sensitivity analysis indicates that extended estimation windows, reduced model orders, and smaller decay factors collectively enhance forecasting precision. Moreover, through comprehensive predictor analysis, we identify commodity market and economic policy uncertainty as significant determinants of the lowest VaR across both 95% and 99% confidence levels. This demonstrates that different predictor impacts depend on market phase and confidence threshold selection.

Leveraging Asynchronous Cross-border Market Data for Improved Day-Ahead Electricity Price Forecasting in European Markets

Presenter: Maria Margarida Mascarenhas (KU Leuven)
Author(s): Maria Margarida Mascarenhas;Jilles De Blauwe;Hussain Kazmi

Accurate short-term electricity price forecasting (EPF) is crucial for strategically scheduling demand and generation bids in European day-ahead markets (DAM), especially amid increasing cross-border market integration and price volatility driven by growing shares of variable renewable energy. This study investigates whether incorporating asynchronously published day-ahead electricity prices from interconnected markets, made possible by differing gate closure times across Europe, improves forecasting accuracy for the Belgian (BE) and Swedish (SE3) bidding zones. We utilize two state-of-the-art forecasting models: a linear, LASSO Estimated Auto-Regression (LEAR) model and a nonlinear, data-driven Deep Neural Network (DNN). Our findings reveal significant improvements in forecast accuracy when including price data from interconnected markets (Germany-Luxembourg, Austria, and Switzerland), due to their strong market coupling and interdependencies. This improvement is evident not only for general market conditions but also for extreme price events, representing both the lowest and highest 5% of observed prices. We additionally explore computational performance by examining the impact of varying calibration window lengths and recalibration frequencies. Our analysis identifies a clear trade-off: frequent recalibration significantly improves accuracy but at substantial computational costs. We therefore provide insights into optimizing these parameters to achieve an effective balance between forecasting accuracy and computational efficiency. Overall, these findings underline the substantial opportunities and inherent challenges associated with leveraging asynchronous market data in electricity price forecasting, providing valuable guidance for market participants and decision-makers aiming to optimize bidding strategies within increasingly interconnected and volatile European energy markets.

Forecasting Considering Dependencies Between Wind Power Generation and Electricity Prices in Great Britain

Presenter: Klimis Stylpnopoulos (University of Glasgow)
Author(s): Klimis Stylpnopoulos;Jethro Browell

This study examines the relationship between forecasts of wind power generation and electricity prices in Great Britain using high-resolution wind speed forecast data from the European Centre for Medium-Range Weather Forecasts. Understanding this relationship is essential for improving the information contained in wind and price forecasts, which are critical for electricity market participants. The analysis seeks to identify key dependencies that can help market participants optimise portfolio balancing strategies and assist operators in maintaining grid stability. To achieve this, we use a two-step methodology. First, we calibrate wind speed forecasts at the level of individual balancing mechanism units, utilising weather forecasts published four times per day up to two days in advance. This calibration accounts for variability based on the location and characteristics of wind farms across GB included in the study. We apply both parametric and non-parametric methods to generate probabilistic forecasts for lead times near the settlement period up to two days ahead, and analyse how forecast updates impact prices. Secondly, we use copulas to model dependencies between wind power output and electricity prices. This dual approach enhances our ability to model complex dependencies between renewable generation and market dynamics. By investigating these dependencies, our findings aim to provide actionable insights for electricity market participants seeking to optimise their trading strategies and manage risk more efficiently. The results are expected to improve grid stability and facilitating better integration of wind energy into electricity markets. Moreover, the insights gained may inform future improvements in forecasting models and contribute to policy discussions surrounding renewable energy integration.

OrderFusion: Encoding Orderbook for Probabilistic Intraday Price Prediction

Presenter: Runyao Yu (Delft University of Technology)
Author(s): Runyao Yu;Yuchen Tao;Fabian Leimgruber;Tara Esterl;Jochen Cremer

Efficient and reliable probabilistic prediction of intraday electricity prices is essential to manage market uncertainties and support robust trading strategies. However, current methods often suffer from parameter inefficiencies, as

they fail to fully exploit the potential of modeling interdependencies between bids and offers in the orderbook, requiring a large number of parameters for representation learning. Furthermore, these methods face the quantile crossing issue, where upper quantiles fall below the lower quantiles, resulting in unreliable probabilistic predictions. To address these two challenges, we propose an encoding method called OrderFusion and design a hierarchical multi-quantile head. The OrderFusion encodes the orderbook into a 2.5D representation, which is processed by a tailored jump cross-attention backbone to capture the interdependencies of bids and offers, enabling parameter-efficient learning. The head sets the median quantile as an anchor and predicts multiple quantiles hierarchically, ensuring reliability by enforcing monotonicity between quantiles through non-negative functions. Extensive experiments and ablation studies are conducted on four price indices: 60-min ID3, 60-min ID1, 15-min ID3, and 15-min ID1 using the German orderbook over three years to ensure a fair evaluation. The results confirm that our design choices improve overall performance, offering a parameter-efficient and reliable solution for probabilistic intraday price prediction.

Can Theories Improve Machines?

Presenter: Yi Cao (Xi'an jiaotong-Liverpool University)
Author(s): Yi Cao;Yinuo Wang;Conghua Wen

Can economic theories improve the performance of machine learning models in option valuation? Current literature emphasises ex-ante approaches, either using theoretical models subsequently corrected by machine learning or transferring economic theories into machine learning models through two-phase frameworks. However, these approaches may falter when theoretical and market distributions diverge significantly, particularly in emerging markets like cryptocurrency options. We introduce a novel mid-ante methodology where the model simultaneously learns economic theories and market information. Given options' market data, we synthesise their theoretical values at adjacent or boundary points in both discrete time and term-structure dimensions using a structural economic model. These synthesised values and market data are strategically distributed across different discrete time points and term structures, avoiding direct overlap, creating an augmented training dataset for our neural network. We implement this approach in cryptocurrency options traded on Deribit, an emerging derivative market known for high volatility and frequent departures from the theoretical prices. We evaluate our theory-augmented model using Black-Scholes and Heston models, demonstrating significantly improved performance for the most actively traded options on BTC and ETH cryptocurrencies. Our evaluation comprises multiple benchmarks: the best-performing transfer-learning model for cryptocurrency options (we conduct thorough empirical studies using eight different configurations of the transfer-learning mechanisms and select the best one for the out-of-sample evaluation), a standard feedforward neural network, a machine-corrected neural network, a theory-only neural network trained exclusively on synthesised data, and traditional Black-Scholes and Heston models. Out-of-sample predictions up to one month ahead demonstrate that our proposed model consistently and substantially outperforms all benchmarks. To validate our findings theoretically, we employ simulated data with varying noise levels. As market prices diverge from theoretical values, transfer-learning models exhibit significant discontinuities in their optimisation landscapes, compromising knowledge transfer. Conversely, our theory-augmented model maintains a smooth optimisation landscape, achieving considerably lower training loss.

A Machine Learning Approach to Detect Accounting Frauds

Presenter: Pietro Perotti (University of Bath)
Author(s): Pietro Perotti;Arman Hassanniakalager;Xinyu Ji;Fanis Tsoligkas

This paper introduces a new fraud detection model to the accounting literature using machine learning (ML). This model, which we refer to as LogitBoost, applies ensemble learning to logistic regressions. We show, using seven performance measures assessing the ability to detect fraud, that our model outperforms the methods based solely on logistic regressions or other ML methods used by prior literature. Further, we identify three key methodological concerns in the application of ML methods in prior research: collinearity among the fraud predictors, the approach to optimize the hyperparameters and the approach to address serial frauds. The research design which we use aims at minimizing the biases related to these issues. We show that these methodological choices can indeed influence the conclusions drawn from ML methods used by prior literature. In additional analyses, we show that LogitBoost outperforms the other models we use as benchmarks in the prediction of fraud in subsequent years. Finally, a fraud

index based on the probabilities derived from LogitBoost exhibits a significantly negative association with future firm valuation.

Generalized Mean Absolute Directional Loss for Machine Learning Models in Algorithmic Investment Strategies

Presenter: Paweł Sakowski (University of Warsaw)
Author(s): Paweł Sakowski;Jakub Michańków;Robert Ślepaczuk

Regardless of the selected asset class and the level of model complexity (Transformer, LSTM, Perceptron/RNN), the GMADL loss function produces superior results than standard MSE-type loss functions and has better numerical properties in the context of optimization than MADL (Michańków et al 2024a). Better results mean the possibility of achieving a higher risk-weighted return based on buy and sell signals built on forecasts generated by the given theoretical model estimated using the GMADL versus MSE or MADL function. In practice, GMADL solves the problem of selecting the most preferable feature in both classification and regression problems, improving the performance of each estimation. What is important, through additional parameterization, GMADL also solves the problem of optimizing investment systems on high-frequency data in such a way that they focus on different strategy variants that contain fewer transactions so that transaction costs do not reduce the effectiveness of a given strategy to zero. Moreover, the implementation leverages state-of-the-art machine learning tools, including frameworks for hyperparameter tuning, architecture testing, and walk-forward optimization, ensuring robust and scalable solutions for real-world algorithmic trading.

Using Self-assessment Data in Automobile Insurance Risk Assessment with Dynamic Modelling

Presenter: Xindi Fang (The University of Melbourne)
Author(s): Xindi Fang;Xueyuan Wu

This study explores the integration of self-risk assessment features in the prior risk assessment for usage-based insurance (UBI), when telematics information is not fully available. We test multiple modeling approaches, including the classical generalized linear model (GLM), stochastic gradient boosting machine (GBM), and neural networks, to derive optimal risk estimates. Our findings demonstrate that self-assessment can improve initial risk estimates when combined with traditional features and, in some cases, could even replace classic features in driving risk assessment. However, we acknowledge the potential for bias and misrepresentation in self-reported risk factors. To address this, we discuss strategies to improve the reliability of questionnaire data and emphasize the critical role of telematics integration. It is definite that self-assessed risk factors cannot replace telematics data in predictive performance. We explore strategies for managing self-assessed risk features once the telematics data are fully accessible, assessing whether they should be retained alongside the telematics information or removed entirely. Although self-assessed risk measures may not offer the same predictive power as telematics, they provide a valuable foundation for UBI, enhancing driver awareness of their own risk profile and fostering greater engagement with existing UBI models.

Quantile-based Modeling of Scale Dynamics in Financial Returns for Value-at-Risk and Expected Shortfall Forecasting

Presenter: Xiaochun Liu (The University of Alabama)
Author(s): Xiaochun Liu;Richard Luger

This paper introduces a new semiparametric approach for forecasting Value-atRisk (VaR) and Expected Shortfall (ES) by modeling the dynamics of the conditional scale of financial returns. The conditional scale, defined as the difference between specific quantiles of the return distribution, is modeled using restricted quantile regression specifications. Focusing on downside market risks, the framework derives VaR forecasts from the left-tail quantiles of the rescaled returns, while ES is approximated as the average of these quantiles over levels below the VaR threshold. This approach provides robust, distribution-free estimates of extreme losses, capturing key features of financial returns such as skewness, heavy tails, and leverage effects. Simulation experiments and real-world analysis demonstrate

that the proposed method outperforms established models, including GARCH and joint VaR-ES conditional quantile approaches, in forecasting accuracy. The application to the daily returns of major international stock indices, which includes the COVID-19 pandemic period, highlights the model's effectiveness in capturing risk dynamics.

LLMs for Financial Time Series Forecasting

Presenter: Stanisław Łaniewski (University of Warsaw)
Author(s): Stanisław Łaniewski;Robert Ślepaczuk

We investigate the use of Large Language Models (LLMs) in daily financial time-series forecasting, focusing on the stock indices (Nasdaq) and volatile assets (Bitcoin). Adapting state-of-the-art frameworks such as Chronos and Time-LLM, we benchmark their performance against established baselines, including ARIMA, exponential smoothing, and various deep learning models. Our results suggest that LLMs often achieve on-par accuracy (and superior with some fine-tuning) for capturing complex patterns in both equity and cryptocurrency data. We further examine interpretability and real-time performance. While LLMs can handle zero-shot forecasts over multiple horizons, they impose higher computational requirements and often struggle when faced with drastic market shifts. By comparing LLM-driven forecasts with conventional AutoML ensembles, we highlight advantages in predictive versatility, but also note potential overfitting risks in high-volatility regimes. Drawing from recent research on hybrid and domain-adaptive networks, we explore whether combining classic statistical methods with LLM-based forecasters enhances robustness for intraday spikes and sudden momentum reversals. Our empirical findings show that, although LLMs can detect subtle patterns in crypto market - where deterministic signals are scarce - these gains may dissipate once transaction costs. Overall, this work provides one of the first comprehensive side-by-side assessments of LLM-centered forecasting pipelines against industry-standard approaches for Nasdaq and Bitcoin. We conclude that while LLM-based models offer promising accuracy and flexible generalization, effective deployment in live trading scenarios demands careful calibration, interpretability measures, and risk monitoring. Our analysis outlines practical guidelines for integrating LLM forecasts with established finance workflows, highlighting both the potential and pitfalls of next-generation language-driven market predictions.

Canary in the Coal Mine: Does Abnormal Communication Signal Stock Price Crash Risk?

Presenter: Zhichong Lyu (西南交通大学)
Author(s): Zhichong Lyu;Feng Ma;Yanjie Song

This paper delves deep into the influence of abnormal communication during IPO roadshows on the future stock price crash risk and its underlying mechanism. The research focuses on the interactive exchanges in IPO roadshows and utilizes the Q&A sessions in IPO roadshows as a key data source, as these sessions can fully reflect the dynamic interaction between potential investors and members of the company's management team. To identify abnormal situations in the interactive Q&A, the research leverages the powerful reasoning ability of DeepSeekR1. This model can not only capture abnormalities in the content and tone of the Q&A but also identify abnormal manifestations in many other aspects. The research designs a three-step prompting strategy for DeepSeek R1 and determines 24 abnormal dimensions in four categories, specifically including the abnormal style of answerers, the abnormal style of questioners, content and compliance abnormalities, and process and governance abnormalities. The research results show that there is a significant positive correlation between abnormal communication in IPO roadshows and the future stock price crash risk. Even after solving the endogeneity problem and completing the robustness test, this conclusion remains robust. Specifically, during the IPO roadshow process, abnormal interactive communication makes it difficult for investors to obtain valuable forward - looking signals, preventing them from forming rational expectations of the stock price, reducing the pricing efficiency, and ultimately increasing the risk of future stock price crash risk. The mechanism analysis further reveals that abnormal communication in IPO roadshows mainly transmits the stock price crash risk through the information transparency channel. In addition, a higher level of analyst attention can, to a certain extent, suppress the excessive rise in stock prices caused by abnormal communication, thereby reducing the risk of future stock price crash risk. This research not only enriches the research results on the influencing factors of stock price crash risk but also provides important references for the formulation of IPO policies and risk prevention and supervision work.

Forecasting VIX Futures Returns: Price Persistence or Random Walk?

Presenter: Yaojie Zhang (Nanjing University of Science and Technology)
Author(s): Yaojie Zhang;Zhikai Zhang

The return predictability of VIX futures tended to follow the convention of financial assets that uses the random walk models including historical average (HA) and no-change (NC) and in the efficient market. In this paper, we propose a parsimonious price autoregression (AR-Price) model, theoretically derived from the price persistence in the VIX futures that is intrinsically similar to the volatility behaviors. The forecast accuracy of the AR-Price is more statistically and economically significant than that of HA and NC, which is further supported by the theoretical analysis and simulation evidence under various settings. The AR-Price model also outperforms the broad predictors concerning the stock market and volatility, and can be extensively applied to forecast the returns of those assets whose prices are highly autocorrelated.

Air Traffic Demand Forecast Reconciliation for Brazil: A Comparison of Traditional and Optimal Reconciliation Techniques for Domestic and International Flight Connections

Presenter: Ulrich Gunter (MODUL University Vienna)
Author(s): Ulrich Gunter;Júlia Brandão Calixto;Lucas Turbay Rangel Calixto;Fernando Luiz Cyrino Oliveira;Maurício Franca Lila

Hierarchical forecasting methods often produce more accurate and reliable forecasts through reconciliation strategies by exploiting the inherent structure of the data. This study introduces a hierarchical framework for forecasting air traffic in Brazil, using data from 2000 to 2019 provided by the National Civil Aviation Agency (ANAC). We apply ETS and SARIMA forecasting methods, along with traditional and optimal hierarchical reconciliation techniques, to improve forecast accuracy across different aggregation levels. Our findings offer a robust tool for anticipating air traffic demand trends and highlight that there is no universally superior forecast reconciliation method. The results underscore the potential of hierarchical forecasting to capture the dynamics of Brazil's air travel market, helping optimize resource allocation, infrastructure planning, and tourism management. A key contribution of this research is that employs the classic reconciliation methods widely used in the literature to Brazilian air traffic demand data for the first time. Additionally, this study bridges a knowledge gap by showing the practical application of reconciliation methods in hierarchical forecasting within the Brazilian context, while also emphasizing the general utility of these approaches for complex tourism markets worldwide.

Hotel Demand Forecasting Combining Historic and Advanced Booking Data

Presenter: Clara Cordeiro (Universidade do Algarve, Portugal)
Author(s): Clara Cordeiro;Nuno António;Sara Lopes

Forecasting demand is vital for the hospitality industry. Until now, forecasts have relied upon historical or advanced booking data. This study combines both data types by creating an ensemble of time series and Machine Learning models. Using data from four hotels, it demonstrates that the proposed hybrid-ensemble method outperforms traditional ones. This result highlights the potential of combining approaches and using multisource data in Revenue Management and offers valuable insights for the hospitality industry.

Forecasting Airbnb occupancy for French Polynesia

Presenter: Bozana Zekan (Modul University Vienna)
Author(s): Bozana Zekan;Mondher Sahli;Vincent Dropsy;Ulrich Gunter

The presence of Airbnb in French Polynesia has grown significantly in recent years. Given the country's limited number of hotels - primarily luxury properties owned or managed by international chains - Airbnb has emerged as a potential driver of more inclusive tourism development. By allowing residents to rent out otherwise unused rooms and properties, Airbnb creates opportunities for local income generation and job creation while minimizing

economic leakage abroad. Among the most visited islands in French Polynesia are Bora Bora, Moorea, and Tahiti. This study aims to forecast monthly Airbnb occupancy on these three islands using AirDNA data from early 2016 to early 2025. To achieve this, various spatial and traditional panel econometric techniques will be applied at the listing level, incorporating key economic covariates such as Airbnb and hotel average daily rates (as own and substitute price measures), source-market-weighted real GDP (as an income measure), and a dummy variable accounting for the COVID-19 pandemic. After splitting the data into training and test sets, the forecast accuracy of these models will be evaluated using classical accuracy metrics and compared to that of standard univariate benchmark models.

Non Bank Financial Operations in Emerging Markets and Consequences for the Financial System

Presenter: Vakhtang Charaia (GRUNI)
Author(s): Vakhtang Charaia

Non-bank financial operations (NBFOs) are rapidly expanding in emerging markets, presenting both opportunities and significant challenges for financial system stability. These operations, encompassing activities like money market funds, securitization vehicles, and various lending platforms, operate outside traditional banking regulations, filling credit gaps and fostering financial innovation. However, their growth introduces systemic risks that demand careful consideration. Emerging markets often face limitations in traditional banking infrastructure, restricting access to credit for small and medium-sized enterprises (SMEs) and individuals. NBFOs can bridge this gap, offering alternative financing solutions and promoting financial inclusion. Their flexibility and agility can stimulate economic growth and diversify funding sources, particularly in sectors underserved by conventional banks. Furthermore, they can attract foreign capital and enhance market liquidity. However, the rapid expansion of NBFOs poses substantial threats to financial stability. The lack of stringent regulatory oversight can lead to excessive leverage, maturity mismatches, and interconnectedness, amplifying systemic risks. The opaque nature of many NBFOs makes it difficult to assess their risk profiles, hindering effective monitoring and regulation. Moreover, the absence of deposit insurance and lender-of-last-resort support can exacerbate the impact of financial shocks, potentially triggering contagion effects. Emerging markets, with their often-weaker regulatory frameworks, are particularly vulnerable to these risks. The potential for regulatory arbitrage, where entities exploit loopholes to circumvent regulations, can undermine the effectiveness of national financial supervision. Additionally, NBFOs can facilitate illicit financial flows, including money laundering and tax evasion, posing challenges for governance and financial integrity.

LLM-Based Knowledge Graph Construction and Application for Government Investment Promotion

Presenter: Chen Han (School of Advanced Interdisciplinary Sciences, University of Chinese Academy of Sciences)
Author(s): Chen Han;Yuanyuan Li;Xijin Tang

Investment promotion serves as a critical governmental strategy to attract external capital and stimulate regional economic growth. While governments have issued extensive policy documents containing valuable information on investment patterns and policy trends, these data resources remain largely underutilized. In this paper, we introduce IPG-LLM, an integrated framework combining domain ontology design, model fine-tuning, and prompt engineering to facilitate information extraction and knowledge graph construction from policy documents. The results demonstrate that the IPG-LLM framework significantly enhances the Deepseek-R1-7B model's performance in named entity recognition and relation triple extraction tasks, achieving comparable results to GPT-4o while exhibiting superior capabilities in low-resource scenarios. Based on the constructed knowledge graph, we further explore the application of the Graph RAG in intelligent question-answering tasks, enabling cross-document correlation analysis of policy elements through entity-relation reasoning and semantic retrieval. This study provides an innovative approach to intelligent analysis and decision support in governmental investment promotion, offering practical insights into the application of knowledge graphs for policy text analysis. The resulting knowledge graph data are publicly available at: https://github.com/hanshenmesen/IPG_LLM_KG./subsection%7BOptimism Bias in Technology Foresight: A Stage-Wise Decomposition and Experimental Investigation of Delphi Applications.}Presenter: Xinyu Jiang (Strathclyde Business School, University of Strathclyde)
Author(s): Xinyu Jiang;George Wright

Delphi, a decision-making technique which systematically aggregates the judgments of a group of experts, is widely applied in various contexts, including technology foresight. However, its heavy reliance on expert opinions makes it susceptible to cognitive biases, such as optimism bias. Specifically, experts may underestimate the time required for the realization of an innovation and may exhibit overconfidence in their judgments. While various studies have demonstrated the presence of optimism bias in Delphi, few have attempted to identify at which stage(s) of the Delphi process this bias may emerge or whether Delphi's structure may inadvertently amplify it. Given this gap, this study decomposes the Delphi method into distinct stages, in order to identify points of conceptual vulnerability to optimism bias. In addition, potential solutions to identified process weaknesses are explored, in order to assess whether Delphi can be modified to further improve forecasting reliability. Our study first analyzes the Delphi process and identifies stages where optimism bias is likely to emerge, based on a review of existing literature. Following this, controlled experimentation tests whether modifications to Delphi characteristics (e.g., statement framing) can reduce optimism bias in expert estimations. Early analysis of the extant literature suggests several possible pathways through which optimism bias may manifest at different stages of Delphi, particularly through statement framing and group feedback dynamics. More importantly, characteristics of standard Delphi applications that are intended to reduce bias may, under specific conditions, actually amplify it. A further set of experimental studies are planned to systematically test proposed interventions, such as adjustments in statement framing, to evaluate their impact on the quality of expert judgment. Our findings are designed to contribute to refining the Delphi method for technology foresight, enhancing its reliability, and thus improving policy making in subsequent research and innovation development priorities and funding allocations.

Decoding Market Signals with Self-adaptive Temporal Modeling in Multi-period Decision-making

Presenter: Zhuomin LIANG (Shenzhen University)
Author(s): Zhuomin Liang;Tongtong Xie;Zelong Yi;Yelin Fu

The volatility of agricultural commodity prices presents substantial challenges for resource allocation and strategic planning in multi-period decision-making within supply chains. However, the complex and non-stationary, non-linear nature of agricultural price dynamics renders these fluctuations difficult to model and predict effectively. To address this challenge, we introduce an AI-driven predictive framework that integrates attention mechanisms with random forests (AM-RF). Leveraging a decade-long observation period (from July 26, 2013, to July 25, 2023), we integrate approximately 21,000 news headlines with macroeconomic indicators, while innovatively incorporating the Chain-of-Thought (CoT) strategy to automatically extract sentiment signals from news content using large language models (LLMs). Our results demonstrate that our AM-RF model significantly outperforms benchmark models in multi-step forecasting across various time horizons, short-term (1 day, 5 days), medium-term (60 days), and long-term (180 days) with superior accuracy and robustness. Furthermore, our model captures differences in the time-series contribution of external environmental factors. Specifically, market sentiment emerges as the most influential factor in short- and medium-term price forecasting, while the economic environment assumes a more prominent role in long-term predictions. In contrast, market fundamentals and price costs exhibit relatively marginal contributions, especially in longer forecasting windows. Our study highlights the potential of AI and LLM in addressing the complexities of dynamic market environments and provides robust evidence to support data-driven decision-making frameworks within agricultural supply chains.

Calendar-based Exploratory Time Series Analysis

Presenter: Mitchell O'hara-Wild (Monash University)
Author(s): Mitchell O'Hara-Wild;Cynthia Huang;Rob Hyndman;Matthew Kay

Effective use of exploratory data analysis for time series helps to identify temporal patterns and accurately specify forecasting models. Exploratory time series analysis describes the process of manipulating and visualizing time series to reveal patterns, structures, and other characteristics. This process often involves summarizing data across calendar-based descriptions of points in time at different granularities, such as time within days, weeks, or years. While several commonly used plots exist for visualizing time series, little work has been done to formalize them into

a unified grammar of temporal graphics. Composable grammatical elements provide the flexibility needed to easily visualize multiple seasonality, cycles, and other complex temporal patterns. These modular elements can be combined to create familiar time series graphics, and remixed to create new informative plots. In this talk, I will demonstrate how the composable time-series manipulation and visualization tools in the mixtime and ggtime R packages can be combined to explore a wide variety of temporal patterns. The mixtime package allows manipulation and analysis across multiple temporal granularities and calendar systems. This enables time-specific operations including temporal aggregation, comparison of time points from different granularities or calendars, and joining time series datasets. The ggtime package extends the ggplot2 ecosystem with new grammar elements and plot helpers for visualizing time series data. These additions leverage calendar structures to visually align time points across different granularities and time zones, warp time to standardize irregular durations, and wrap time into compact calendar layouts. These tools provide a highly flexible system for exploring the complex temporal dynamics in modern time series datasets.

Gold Price Forecasting based on CEEMDAN-Conditional Autoencoder Method

Presenter: Yan Yu (Academy of Mathematics and Systems Science, Chinese Academy of Sciences)
Author(s): Yan Yu;Qin Bao

This study proposes a novel decomposition-reconstruction-prediction-integration framework for gold price forecasting based on the Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) and Conditional Autoencoder (CAE) method. Firstly, the CEEMDAN method is used to decompose the gold price series into different modes based on the intrinsic time-scale features. Secondly, according to the fine-to-coarse approach, the different modes are reconstructed as three components, including the long-term trend, the medium-term influence, and the short-term disturbance. Thirdly, the CAE method is used to extract information from candidate explanatory variables and make predictions for the three components. To select the explanatory variables, we incorporate gold's multifaceted attributes, and especially take China's factor and global uncertainty index into consideration. We choose candidate variables by calculating the correlations, and separate them between internal attributes and external environment. The internal attribute variables are used to estimate the common factors through the CAE's factor network, while the external environment features are used to estimate the factor load through the CAE's factor load network. Fourthly, the three parts are integrated to predict the gold price movement as a whole. The results indicate that the CEEMDAN-CAE model proposed greatly improve the prediction for the gold price. On one hand, the decomposition helps reduce the prediction error, demonstrating that the "decomposition-integration" approach works for the prediction of gold price, as modeling distinct time-scale characteristics significantly improves accuracy. On the other hand, the results implies that the CAE method well applies for gold price prediction, and introducing covariates plays more important role than including nonlinearity. Following the contribution calculation, the Geopolitical Risk Index and 10-year TIPS rate are identified as dominant predictors and Chinese factor and global uncertainty index also play an important role in predicting the London Bullion Market Association (LBMA) gold price. This study demonstrates that the asset pricing factor model can be extended and applied to the prediction and analysis of gold price, and the CAE method is effective in information extraction. This study fills a void in the application of factor analysis models in the field of gold price forecasting.

Mortality Modelling Using Linked Death Registration and Census Data

Presenter: Han Li (The University of Melbourne)
Author(s): Han Li

Recent availability of death registration data linked at the individual level to the census data provides enormous potential to improve measurement of inequalities in mortality. Machine learning techniques including random survival forests, survival support vector machines, and extreme gradient boosting, have been recognised as powerful tools in survival analysis. In this research, we propose an explainable machine learning approach to identify the drivers of mortality across different causes of deaths.

Metabolomic Machine Learning Predictor for Diagnosis and Prognosis of Gastric Cancer

Presenter: Yizi Zhao (Tsinghua university)
Author(s): Yizi Zhao;Bohong Wang;Yangzi Chen

Gastric cancer (GC) represents a significant burden of cancer-related mortality worldwide, underscoring an urgent need for the development of early detection strategies and precise postoperative interventions. However, the identification of non-invasive biomarkers for early diagnosis and patient risk stratification remains underexplored. Here, we conduct a targeted metabolomics analysis of 702 plasma samples from multi-center participants to elucidate the GC metabolic reprogramming. Our machine learning analysis reveals a 10-metabolite GC diagnostic model, which is validated in an external test set with a sensitivity of 0.905, outperforming conventional methods leveraging cancer protein markers (sensitivity < 0.40). Additionally, our machine learning-derived prognostic model demonstrates superior performance to traditional models utilizing clinical parameters and effectively stratifies patients into different risk groups to guide precision interventions. Collectively, our findings reveal the metabolic landscape of GC and identify two distinct biomarker panels that enable early detection and prognosis prediction respectively, thus facilitating precision medicine in GC

Myopia and High Myopia Trends in Chinese Children and Adolescents Over 25 Years: A Nationwide Study with Projections to 2050

Presenter: Zhe Pan (Beijing Tsinghua Changgung Hospital)
Author(s): Zhe Pan;Feng Li;Ya Xing Wang;Tien Yin Wong

Background: The global rise in myopia, particularly in Asia, presents significant public health challenges. Analyzing trends and forecasting impacts are critical for developing strategies to mitigate this burden. Methods: We conducted the largest study to date on myopia and high myopia prevalence in Chinese children and adolescents aged 7–18 years, analyzing data from 5,095,256 individuals across 119 studies from 1998 to 2022. Data variability between cycloplegia and non-cycloplegia measurements was addressed using a distance-based model averaging calibration. Aggregated prevalence and age-specific urban-rural trends were estimated using thin plate spline regression, with projections to 2050 derived from time series modeling. Findings: Myopia prevalence plateaued in 2006 in urban areas and in 2013 in rural areas, with the urban-rural prevalence gap narrowing since 2015 (urban/rural ratio below 1·3 for all ages). By 2050, myopia prevalence is projected to stabilize at 27·1% (95% CI: 10·0–44·4%) for ages 7–9 and 81·5% (74·7–88·3%) for ages 16–18 in urban areas, and at 20·1% (8·6–31·7%) and 74·1% (63·2–84·8%), respectively, in rural areas. High myopia prevalence among adolescents aged 16–18 is expected to rise from 7·3% in 2001 to 22·1% by 2050. Prevalence correlated significantly with the Human Development Index ($P < 0·001$). Interpretation: Despite stabilization in overall myopia prevalence, the continued rise in high myopia underscores the need for targeted control measures. Projections emphasize the importance of addressing regional disparities and prioritizing public health interventions.

Loss-based Bayesian Sequential Prediction of Value-at-Risk with a Long-Memory and Non-linear Realized Volatility Model

Presenter: Chao Wang (The University of Sydney)
Author(s): Chao Wang;Rangika Peiris;Minh-Ngoc Tran;Richard Gerlach

A long-memory and non-linear realized volatility model class is proposed for direct Value-at-Risk (VaR) forecasting. This model, referred to as RNN-HAR, extends the heterogeneous autoregressive (HAR) model, a framework known for efficiently capturing long memory in realized measures, by integrating a Recurrent Neural Network (RNN) to handle the non-linear dynamics. Quantile loss-based generalized Bayesian inference with Sequential Monte Carlo is employed for model estimation and sequential prediction in RNN-HAR. The empirical analysis is conducted using daily closing prices and realized measures with around 12 years of data till 2022, covering 31 market indices. The proposed model's one-step-ahead VaR forecasting performance is compared against a basic HAR model and its extensions. The results demonstrate that the proposed RNN-HAR model consistently outperforms all other models considered in the study.

Forecasting of the Korean Stock Market Volatility Index Using a Modified Multi-Input LSTM Model

Presenter: Heejoon Han (Sungkyunkwan University)
Author(s): Heejoon Han;Gyure Kim

This study explores the forecasting of the VKOSPI index, a vital volatility indicator for the Korean stock market, using advanced machine learning models. VKOSPI, similar to the VIX index in the U.S., reflects market expectations of future volatility and is crucial for risk management and financial derivatives pricing. Traditional forecasting models, such as the heterogeneous autoregressive (HAR) model, have been widely used in this domain. However, this research introduces a modified Multi-Input Long Short-Term Memory (LSTM) model, which incorporates a wide range of explanatory variables, including general financial variables, stock market fund flows, and stock market transaction trends, to improve predictive accuracy. Our Multi-Input LSTM model builds upon and expands on the work of Li et al.(2023), featuring a sequential layer aggregation design that enables it to capture complex patterns often overlooked by simpler models. The study, based on data from January 2016 to March 2023, compares the forecasting performance of the Multi-Input LSTM model with various other machine learning models, such as random forest, XGBoost, and linear models like LASSO and Elastic Net. With a total of 216 explanatory variables, including lagged values, the models are tested across different forecast horizons, including 1 day, 5 days, 10 days, and 22 days. The results show that machine learning models, particularly non-linear ones such as LSTM, Random Forest, and XGBoost, outperform the HAR model, especially for longer forecast periods. The Multi-Input LSTM demonstrates superior accuracy, particularly when variables related to stock market fund flows and stock market transaction trends are included, highlighting their importance in predicting volatility. Furthermore, the study uses rigorous statistical tests to validate the enhanced performance of the Multi-Input LSTM model over traditional approaches. The study contributes to the growing field of machine learning in finance, demonstrating that advanced models like the Multi-Input LSTM can effectively handle the complexities of financial time series data.

Neural Network Sieves for Semiparametric Copula Estimation

Presenter: Hasan Fallahgoul (Monash University)
Author(s): Hasan Fallahgoul;Artem Prokhorov

We propose a neural network sieve approach for semiparametric copula estimation, adapting scale-insensitive deep neural networks to approximate the unknown copula density. The method extends existing Bernstein-Kantorovich polynomial sieves with a more flexible neural network architecture while maintaining theoretical guarantees. We establish consistency, asymptotic normality, and semiparametric efficiency of the sieve maximum likelihood estimator for the marginal distribution parameters. The approach addresses computational and theoretical challenges through scale-insensitive network constraints and careful normalization to preserve copula properties.

Market Timing with Bi-Objective Cost-Sensitive Machine Learning

Presenter: Artem Prokhorov (U Sydney, CEBA, CIREQ)
Author(s): Artem Prokhorov;Robert James;Jessica Leung

This paper develops a framework for cost-sensitive training of machine learning models that predict the direction of aggregate stock returns. We design a bi-objective loss function that augments the traditional log-loss objective with an objective that minimizes the cost of indi- vidual false-positive and false-negative classification errors. We argue that the option-implied conditional value-at-risk is a natural measure of the misclassification costs in such models. Our bi-objective optimization framework permits us to isolate the effect of cost-sensitivity from log-loss minimization, and to integrate forward-looking information from options markets di- rectly into the model training process. We study changes in the classification performance of elastic-net logistic regression and gradient boosted decision trees trained using the bi-objective framework. The new approach improves the risk-adjusted returns of market timing strategies and substantially reduces downside risk.

Detection of Explosive Bubbles in Time Series Using Mixed Integer Programming

Presenter: Alexander Semenov (University of Florida)
Author(s): Alexander Semenov;Artem Prokhorov;Anton Skrobotov;Peter Radchenko

Identifying rational bubbles in time series data is a key focus in econometrics and applied statistics due to their significant impact on banking systems and macroeconomic stability. While traditional approaches have primarily relied on statistical methods, recent research highlights the potential of Mixed Integer Programming (MIP) techniques to address complex statistical challenges. In this paper, we propose a MIP-based method specifically designed to identify and estimate bubbles in time series data. Our approach extends the structural break detection framework by formulating bubble detection as a Mixed Integer Quadratic Programming problem. The framework is built upon a least squares formulation, incorporating a penalty on both the breakpoints and the bubble regimes. We present the complete MIP formulation and evaluate its performance using various real and synthetic datasets.

Research on Spatiotemporal Mixed-Frequency Data-Driven Nonlinear Forecasting Model for Provincial GDP in China: A Machine Learning Perspective

Presenter: Ni Zhang (Chongqing Normal University)
Author(s): Ni Zhang;Xianning Wang

Against the backdrop of regional economic coordinated development and accelerated spatiotemporal flows of economic factors, regional economic growth exhibits significant spatiotemporal evolution characteristics. Direct utilization of information carried by high-frequency/low-frequency data is crucial for improving the accuracy and timeliness of regional GDP forecasting. Considering the spatiotemporal interaction effects of GDP and mixed-frequency data features, this study constructs a spatiotemporal interaction autoregressive mixed-data sampling support vector regression model (S-AR-Midas-SVR) incorporating a spatial geographical weight matrix. The model is applied to annual and quarterly GDP data and economic indicators of 30 Chinese provinces (excluding Hong Kong, Macao, Taiwan, and Tibet) from 2010 to 2023, enabling spatiotemporal mixed-frequency forecasting analysis under the interaction of GDP temporal lags and spatial lags. Empirical results demonstrate that: 1) The model achieves the smallest prediction error and best fitting performance among nine forecasting models, including contemporaneous mixed-frequency autoregressive and temporal lag mixed-frequency autoregressive models. A training-test data partition ratio of 8:2 yields superior forecasting results in cross-validation; 2) In provincial forecasting accuracy comparisons, the model shows strong explanatory power, with an overall mean squared error (MSE) controlled within 0.2. The MSE values for 22 provinces/municipalities (e.g., Henan, Xinjiang) are less than 0.1; 3) Cross-regional heterogeneity tests reveal that due to economic development levels and spatial agglomeration effects, forecasting accuracy decreases in a gradient pattern from eastern to central and western regions. The proposed spatiotemporal mixed-frequency forecasting framework provides a novel method for regional economic monitoring and quantifies the basis for cross-regional economic policy coordination.

Finance-Informed Neural Networks

Presenter: Zhibin Deng (University of Chinese Academy of Sciences)
Author(s): Zhibin Deng;Hongxu Wu

Inspired by Physics-Informed Neural Networks (PINN), we introduce a novel framework called Finance-Informed Neural Networks (FINN), which integrates financial theory with deep learning technology to enhance the transparency and accuracy of deep learning methods in empirical financial research. The network structure of FINN is based on the Arbitrage Pricing Theory (APT) and empirical portfolio construction techniques, with market efficiency information incorporated into the loss function during training. Empirical studies on China's A-share market demonstrate that FINN offers several advantages. First, in terms of out-of-sample predictive performance, FINN outperforms conventional fully connected neural networks and traditional APT factor models, achieving higher R-squared values. Second, FINN overcomes the limitations of traditional APT factors in predicting different asset types, showcasing stronger generalization capabilities. Additionally, when constructing mean-variance efficient portfolios, FINN achieves the highest cumulative returns and Sharpe ratios, highlighting its significant economic value.

Research on Unmanned Warehouse Demand Interval Forecasting Based on Ensemble Kernel-free Optimal Fuzzy Margin Distribution

Presenter: Yukai Zheng (International Business School, Hainan University, Haikou, 570228, China)
Author(s): Yukai Zheng;Jian Luo

In response to the challenges of uncertainty and suddenness in order demands within unmanned warehouses, this study proposes an interval forecasting method based on kernel-free optimal fuzzy margin distribution quantile regression averaging. First, a well-designed fuzzy membership function is introduced to assign differentiated weights to samples, thereby effectively mitigating the impact of outliers on the model. Second, the concept of optimal margin distribution is incorporated into the framework of kernel-free support vector regression. By optimizing the margin distribution of the samples, the robustness and generalization capability of the model are significantly enhanced. Finally, in combination with the principles of ensemble learning, multiple feature selection methods are employed to construct variable subsets for model training. The integration of multiple models is achieved via quantile regression averaging, which not only improves the accuracy and efficiency of interval forecasting but also ensures stability. Extensive experimental results based on real-world data demonstrate that the proposed method exhibits outstanding robustness and superiority in predicting order demands in unmanned warehouses, thus providing reliable scientific support for inventory management and personnel scheduling in complex dynamic environments.

Human-Centric Order Picking: Performance Prediction and Robot Assignment at a Robotic Fulfilment Center

Presenter: Wu Zhiqiao (Dongbei University of Finance and Economics)
Author(s): Wu Zhiqiao;Luo Jian

Abstract: Problem definition: E-commerce giants scale up their order-picking operations by adopting robotic fulfill-ment centers (RFCs). In RFCs, automated guided vehicles (AGVs) transport movable shelf racks to pickers' workstations, instead of having human pickers travel to pick items. Unfortunately, this apparent relief for pickers turns out to be a curse: Pickers become the bottleneck in the order-picking process. They suffer from high-intensity, stationary, and repetitive work, and consequently, from both physical and mental health problems. To ease this tension, we collaborate with a major e-commerce firm to study how the RFC can improve picking efficiency by considering heterogeneous picker performance. Methodology/results: We first propose a distributionally robust human-centric picking performance prediction (DHPP) model to predict two performance metrics of pickers: picking time and performance inconsistency. We incorporate probabilistic constraints to address the imperfect information about the random features associated with the picker's behavior. We show that the DHPP model can be reformulated into a tractable second-order cone program. Based on the predicted performance metrics, we then propose a mixed 0–1 program to optimize the picker and order assignments. Managerial implications: Our computational study demonstrates that the DHPP model significantly outperforms the state-of-the-art forecasting models regarding prediction accuracy. Our simulation, calibrated with real data from JD.com, shows that our strategy reduces the number of remaining unfulfilled items by 14.2% and improves average pickers' picking productivity by 7.5%. These results suggest that our strategy can significantly improve pickers' welfare by increasing their income and alleviating their mental stress and illness.
Keyword: Robotic fulfillment centers; order picking; distributionally robust performance prediction; human-centric order assignment

Multi-view Learning with Enhanced Multi-weight Vector Projection Support Vector Machine

Presenter: Xin Yan (Shanghai University of International Business and Economics)
Author(s): Xin Yan;Shuaixing Wang;Huina Chen;Hongmiao Zhu

Multi-view learning aims on learning from the data represented by multiple distinct feature sets. Various multi-view support vector machine methods have been successfully applied to classification tasks. However, the existed methods often face the problems of long processing time or weak generalization on some complex datasets. In this paper, two multi-view enhanced multi-weight vector projection support vector machine models are proposed. One is a ratio form of multi-view enhanced multi-weight vector projection support vector machine (R-MvEMV),

while the other is a difference form (D-MvEMV). Instead of searching for specific classification hyperplanes, each proposed model tries to generate two projection matrices composed of a set of projection vectors for each view. A co-regularization term is added to maximize the consistency of different views. R-MvEMV and D-MvEMV can be simplified to two generalized eigenvalue problems and two eigenvalue problems, respectively. The optimal weight vector projections are the eigenvectors corresponding to the smallest eigenvalues. Some numerical tests are conducted to compare the proposed methods with the other state-of-art multi-view support vector machine methods. The numerical results show the better classification performance and higher efficiency of the proposed methods.

Probabilistic and Point Forecast of Daily Sales in Walmart Stores

Presenter: Slawek Smyl (Walmart Labs)
Author(s): Slawek Smyl;Fiona Yeung;Johann Posch

Forecasting retail sales is challenging, due to complicated seasonality, impact of holidays and events, weather, inventory, pricing, cannibalization, sparseness etc. Typically, it is done at weekly intervals, and doing it on daily data, increases the difficulty. We will report on a new 28 days-ahead system that outputs both point forecast (expected values) and 10 quantiles, so it also approximates a probabilistic forecast. The probabilistic part operates on summed values along time, so e.g. x3:=x1+x2+x3. This aligns better with the expected use as an input to an ordering optimization system. We will also describe some challenges and solutions in productionalization of the system at Walmart scale.

Probabilistic Forecasts with Global Gradient Boosted Decision Trees: Rethinking Bootstrapping with Structural Adjustments

Presenter: Filotas Theodosiou (VIVES University of Applied Sciences)
Author(s): Filotas Theodosiou;Yves R. Sagaert;Liselot De Vlieger

Accurate probabilistic forecasts are vital for large-scale operational tasks such as inventory management, where decision-making under uncertainty depends on quantifying risks effectively. Although global Gradient Boosted Decision Trees (GBDTs) have gained popularity for their strong point forecast accuracy in such settings, their potential for reliable prediction intervals remains underexplored. Existing prediction interval estimators — such as quantile regression, empirical methods, and conformal predictions — often do not fully leverage GBDTs' advantages. This paper argues that bootstrapping is better suited for global GBDTs as it aligns with their autoregressive nature and internal mechanisms, providing local corrections for patterns overlooked during global optimization. We propose enhancements to conventional bootstrapping to address limitations like neglected temporal dependencies and non-representative error sampling. Specifically, we refine bootstrapping with structured residual sampling, validation-based error selection, magnitude-specific binning, and kernel density estimation for quantile extraction. We evaluate these modifications on three demand forecasting case studies in the food retail sector in Flanders, Belgium, and demonstrate that our modified bootstrapping consistently outperforms established baselines in prediction interval accuracy and downstream inventory performance. Additionally, we provide recommendations for specific modifications under different dataset characteristics and analyze the correlation between point forecast and prediction interval errors. We further show that our methods do not have constraining computational requirements. The significance of our findings lies in the enhancement of GBDTs, a forecasting tool adopted widely by practitioners, with a flexible prediction interval estimator that efficiently improves predictive reliability across diverse scenarios.

Enhancing the Utility of Forecasts through Decision-Informed Forecasting

Presenter: Yiwen Wang (The University of Queensland)
Author(s): Yiwen Wang;Abolghasemi Mahdi

Prediction + Optimization problems are widely studied in academia and industry. However, existing research shows that improvements in forecasting accuracy do not necessarily lead to optimal decision-making outcomes. To reduce this discrepancy, we propose Decision-Informed Forecasting (DIF), a framework that integrates decision values into training predictive models. We apply DIF to hyperparameter optimization by incorporating decision

performance as an objective in the tuning process. In addition, we introduce a multi-objective optimization strategy that optimizes both decision quality and predictive accuracy to refine hyperparameter selection further. We evaluate the approach on M4 Industry dataset and evaluate inventory decisions using Newsvendor and the Order-Up-To Policy. Experimental results demonstrate that single-objective and multi-objective DIF frameworks outperform traditional two-step methods on final decision-making performance.

Learning Data-Driven Uncertainty Set Partitions for Robust and Adaptive Energy Forecasting with Missing Data

Presenter: Akylas Stratigakos (Imperial College London)
Author(s): Akylas Stratigakos;Panagiotis Andrianesis

Short-term forecasting models typically assume input data (features) availability when deployed and in use. However, equipment failures, disruptions, and cyberattacks may lead to missing features when such models are used operationally, which could negatively affect forecast accuracy, and result in suboptimal operational decisions. In this paper, we use adaptive robust optimization and adversarial machine learning to develop forecasting models that seamlessly handle missing data operationally. We propose linear- and neural network-based forecasting models with parameters that adapt to available features, combining linear adaptation with a novel algorithm for learning data-driven uncertainty set partitions. The proposed adaptive models do not rely on identifying historical missing data patterns and are suitable for real-time operations under stringent time constraints. Extensive numerical experiments on short-term wind power forecasting considering horizons from 15 minutes to 4 hours ahead illustrate that our proposed adaptive models are on par with imputation when data are missing for very short periods (e.g., when only the latest measurement is missing) whereas they significantly outperform imputation when data are missing for longer periods. We further provide insights by showcasing how linear adaptation and data-driven partitions (even with a few subsets) approach the performance of the optimal, yet impractical, method of retraining for every possible realization of missing data.

A Novel Robust Ensemble Learning Method for Load Forecasting under Large-scale Data Attacks

Presenter: Xuqiang Liu (Hainan University)
Author(s): Xuqiang Liu;Jian Luo;Yukai Zheng

Abstract Load forecasting is significantly crucial for optimal operational management of smart grid within the energy industry. As power systems undergo accelerated digitization and face increasingly frequent and sophisticated data attacks, it is highly necessary to develop robust methodologies for producing timely and accurate load forecasts under data attacks. This is critical to ensuring grid stability, optimizing resource allocation, and mitigating the cascading effects of potential cyber-physical disruptions. In this study, a novel ensemble quadratic framework based iteratively reweighted L1 (EnQF-IRL1) method is proposed for load forecasting under large-scale data attacks. We first propose the iteratively reweighted L1 (IRL1) model by integrating the L1-norm loss function with adaptive reweighting techniques for robust load forecasting under data attacks. Then a weight function based on IRL1 is introduced for efficiently calculating the relative importance of each observation in the load history. Third, we propose a novel ensemble quadratic framework (EnQF) by incorporating the weight function and L1-norm regularization for enhancing the forecasting performance. Finally, we incorporate the IRL1 model and ensemble quadratic framework to propose a novel robust EnQF-IRL1 method by incorporating all iterative forecasting outputs of the IRL1 into original data features for final load forecasting. Comprehensive numerical results on benchmark dataset GEFCom 2012 demonstrate the superior performance of proposed En-IRL1-QR method over well-known robust load forecasting models and ensemble learning methods without or with two types of data attacks. Keywords: Load forecasting, Data attacks, Weight function, Ensemble learning, Quadratic regression model, Ensemble quadratic framework(EnQF)

Prediction of Chemical Oxygen Demand (COD) in Wastewater Effluent: A Hybrid Approach Combining VMD, Bi-LSTM, and Genetic Algorithm Optimization.

Presenter: Fang Wang (School of Economics & Management, Xidian University)
Author(s): Fang Wang;Qingcheng Hu;Lei Hong;Guoqing Yang

Abstract: The concentration of Chemical Oxygen Demand (COD) in the effluent serves as a crucial indicator for evaluating the performance of wastewater treatment. Accurately predicting the COD concentration in the effluent of wastewater treatment plants is highly valuable for optimizing energy consumption, controlling operating costs, and enhancing treatment efficiency. Due to the nonlinearity, frequency fluctuations, and complexity of the COD time-series data, a parameter optimization method based on Bayesian optimization and the Genetic Algorithm (GA) is developed. A COD concentration prediction method is proposed, combining Variational Mode Decomposition (VMD) and Bidirectional Long Short-Term Memory (Bi-LSTM) networks. First, the Spearman correlation coefficient method is utilized to select indicators that significantly correlate with the effluent COD concentration. Next, a Genetic Algorithm is employed to optimize the number of VMD decomposition components, which involves decomposing the original COD time series into multiple stationary Intrinsic Mode Functions (IMF) components. Following this, Bayesian optimization is applied to fine-tune the hyperparameters of Bi-LSTM to predict each IMF component. Ultimately, the predictions of each IMF component are integrated to yield the final COD concentration forecast. The modeling results based on actual data from a wastewater treatment plant in Zhaoqing, Guangdong Province, China, show that the proposed method has a significant advantage in prediction accuracy compared to traditional models such as LSTM and VMD-LSTM. The average absolute percentage error (MAPE) and the coefficient of determination (R^2) are 0.80% and 0.96, respectively, indicating a notable improvement in the accuracy of COD concentration prediction in wastewater treatment plant effluent. Keywords: wastewater treatment; Variational Mode Decomposition; Bidirectional Long Short-Term Memory network; prediction

Construction and Application of China's Weekly Economic Index (WEI): Real-Time Monitoring and Asset Allocation Based on High-Frequency Data

Presenter: Yun Bai (E Fund Management Co., Ltd.)
Author(s): yun bai;Xian Han

This study develops a novel Weekly Economic Index (WEI) for China to overcome the limitations of traditional low-frequency indicators like GDP that fail to provide timely economic signals. Building upon the U.S. WEI framework (Lewis et al., 2020), we construct a comprehensive high-frequency monitoring system using non-parametric methods and innovative data processing techniques. Our approach differs fundamentally from conventional nowcasting models by focusing on continuous weekly economic dynamics rather than predicting specific low-frequency targets. The research yields four key contributions: First, we demonstrate that our standardized WEI maintains strong correlations ($R^2 \geq 85\%$) with both GDP and PMI while capturing finer intra-month economic fluctuations. Second, we develop and validate a specialized seasonal adjustment methodology tailored for high-frequency economic data. Third, empirical results show the WEI's superior performance in asset allocation, with timing strategies based on our index consistently outperforming traditional momentum approaches across multiple asset classes. Fourth, we establish a robust framework for handling challenging data characteristics including small samples, mixed frequencies, and missing observations. The practical implications are significant: policymakers gain a powerful tool for real-time economic monitoring and policy evaluation, while investors benefit from improved market timing signals. Although the study acknowledges potential limitations regarding indicator persistence, the modular architecture ensures efficient updates. Future research directions include expanding indicator coverage and enhancing sectoral balance to further improve the index's robustness. Keywords: High-frequency economic indicators, Weekly Economic Index, Real-time monitoring, Seasonal adjustment, Asset allocation, Economic nowcasting

Nowcasting and Forecasting China's Industrial Value-added Growth based on an Ensemble Approach

Presenter: Qin Bao (Center for Forecasting Science, Academy of Mathematics and Systems Science, Chinese Academy of Sciences)

Author(s): Qin Bao;Yan Yu;Yangyang Zheng

Industrial economy is the pivot of China's economy, and accurate industrial growth prediction would provide effective policy support. Nowadays, there are multiple high-frequency indicators related with industrial growth, such as daily and weekly data for operating rate or product output, which contain rich information for early prediction. By utilizing mixed-frequency data, we build the nowcasting and forecasting model for China's monthly industrial value-added growth based on an ensemble approach. Firstly, based on the variable selection from a large group of mixed-frequency data, both linear and nonlinear models are developed with univariable and multivariables. Secondly, by using the bagging and boosting ensemble algorithms, the results are integrated and evaluated based on forecasting accuracy and directional consistency. For the univariate models, we use the classic linear time series forecasting model, the Autoregressive Integrated Moving Average (ARIMA) model and the nonlinear Long Short-Term Memory (LSTM) model. For the multivariate models, we select the classic linear Vector Autoregressive (VAR) model, mixed-frequency MIDAS and MFDFM models, and the nonlinear Random Forest (RF) model. For the mixed-frequency models, the explanatory variables are selected based on LASSO, Ridge, and Elastic Net method. The nowcasting results indicate that among the univariate models, the ARIMA model with parameters selected based on the RMSE minimization rule achieves higher prediction accuracy. Moreover, for the multivariate models, only the RF model provides better prediction than the ARIMA model, implying high nonlinearity amongst industrial related variables. The multi-model ensemble approach offers relatively higher prediction accuracy, with the Boosting ensemble algorithm yielding the best results. Long-term forecasting indicate that the RF model has the advantage over others. This research shed light on the prediction of China's industrial value-added growth.

From Micro to Macro: Can Machines Unlock New Real-Time Economic Signals from Firm-Level Accounting Data?

Presenter: Shijie Zhu (Central University of Finance and Economics)
Author(s): Shijie Zhu;Naijing Huang

We apply machine learning (ML) techniques to uncover new real-time economic signals from firm-level accounting data for weekly updated nowcasts of U.S. economic output. Accurate and timely economic assessments are crucial for policymakers, investors, and businesses, yet official Gross Domestic Product (GDP) figures are typically released with significant delays. Our approach addresses this gap by employing multiple ML models to utilize firm-level accounting data, demonstrating substantial improvements in nowcasting accuracy over a random walk benchmark. To validate the robustness of our methodology, we assess its predictive stability across varying economic conditions, including high-volatility periods, and evaluate its performance using alternative output measures, such as nominal GDP and Gross Domestic Income (GDI). The consistent performance across these tests underscores the reliability of our ML-driven approach. By decomposing GDP into its major components and sectoral gross value added, we identify strong predictive signals from firm-level data for key GDP subcomponents, including personal consumption, government spending, and non-residential investment, while net exports remain less predictable. Nowcasts of sector-level gross value added growth reveal particularly strong predictability within the business sector, underscoring the close link between accounting data and business-sector dynamics. We further compare our framework with alternative modeling approaches. Our ML-based approach outperforms both models using aggregate accounting series and dynamic factor models. Moreover, by integrating ML techniques with mixed data sampling (MIDAS) regressions, we demonstrate that our approach not only surpasses models that rely solely on mixed-frequency macroeconomic variables but also that incorporating macro and financial variables into the accounting data does not yield additional predictive gains. Finally, a variable importance analysis identifies firms in the Industrials and Consumer Discretionary sectors as particularly influential. A subsequent decomposition at the variable level further emphasizes profitability metrics as crucial indicators of economic activity. Overall, our study demonstrates that directly analyzing firm-level accounting data through ML methods substantially advances macroeconomic nowcasting, offering critical insights for proactive economic decision-making and real-time macroeconomic monitoring.

Nowcasting U.S. Hog and Pig Inventory

Presenter: Tao Xiong (Huazhong Agricultural University)
Author(s): Tao Xiong;Lee Schulz;Xinyue He;Dermot Hayes;Wendong Zhang

U.S. hog and pig inventory data are one of six principle economic indicators of the U.S. agricultural economy published by the National Agriculture Statistics Service (NASS). This data is published on a quarterly basis. During times of economic disruption in the pork value chain such as the initial outbreak of porcine epidemic diarrhea virus and the COVID-19 pandemic, there is a need for more frequent reporting of this inventory data. These more frequent data, if accurate, would improve production and marketing decisions, aid in policy assessments, and help reduce price volatility. This study proposes a dynamic factor model (DFM) to nowcast inventory values published in NASS quarterly Hogs and Pigs reports from 1993 to 2024 using more frequent production data from the U.S. Department of Agriculture and futures price data from the Chicago Mercantile Exchange. In each round of nowcasting between two consecutive quarterly Hogs and Pigs reports, we obtain ten nowcasting predictions of hog and pig inventory. Our results show that the nowcasting model yields accurate predictions in the months and weeks ahead of the release of the next Hogs and Pigs report. The DFM model outperforms several counterparts, especially during extreme events. Government agencies and professional organizations can potentially utilize our nowcasting model to publish inventory estimates more frequently and at lower costs.

A Parallel NSGA-III with Solution Space Compression for MDVRPTW in Refined Oil Distribution with Multiple Scenarios

Presenter: Qi Peng (China University of Petroleum (East China))
Author(s): Peng Qi;Xiaofeng Xu

As a crucial energy resource for national livelihood and market stability, refined oil involves complex business scenarios and market demands. Its distribution process strongly emphasizes model generalizability and algorithmic computational efficiency. This study investigates the detailed operational logic and scheduling rules of refined oil distribution, establishing a multi-depot vehicle routing optimization model with time windows that balances cost minimization, satisfaction maximization, and loading rate maximization. To enhance model adaptability across different scenarios, modular modeling strategies are incorporated for applicability improvements. Additionally, considering the requirement for real-time optimization, this study designs a “prediction + optimization” solution approach and develops a parallel NSGA-III with solution space compression (PNSGA-III-SSC) to improve overall solving efficiency. Finally, empirical validation is conducted using real-case data from multiple scenarios of varying scales. The results demonstrate that compared with actual distribution solutions, the proposed model better meets the demands of refined oil marketing enterprises and significantly enhances distribution efficiency. The PNSGA-III-SSC also demonstrates superior performance compared to benchmark algorithms.

Uncovering Retailers′ Inventory Policies: A Data-Driven Approach Using Historical Order Data

Presenter: None
Author(s): Mostafa Rafienezhad Masouleh;Kai Hoberg

Managing inventory effectively is a critical challenge for retailers, as both excessive and insufficient stock levels result in financial and operational inefficiencies. Accordingly, operations management research has extensively investigated optimal inventory policies under different assumptions. Many of these policies have been embedded into ERP systems, which are now widely used in practice. However, despite the central role of inventory policies in order fulfillment, little research has been conducted that aims to understand which of these policies are applied by retailers, which could have important implications for demand planning of suppliers. This study aims to bridge this gap by systematically analyzing real-world inventory policies using a large empirical order dataset. We utilize data from a leading global wholesaler in the travel retail market, which supplies hundreds of retailer customers across diverse segments, including airports, ships, and border shops. The dataset comprises order histories for tens of thousands of SKUs across multiple product categories, spanning two years. The study’s primary objectives are: (1) to characterize the inventory policies of the retailers based on observed order patterns and (2) to identify deviations from standard policies to understand the factors influencing non-standard ordering behaviors. To achieve these objectives, the study employs a machine learning-based approach, integrating Recurrent Neural Networks. By this approach, the model processes raw sequential order data directly, enabling it to capture demand fluctuations and

replenishment patterns over time. Preliminary findings indicate that real-world inventory policies frequently deviate from theoretically optimal models due to various operational and behavioral constraints. Retailers often override automated replenishment systems in response to unforeseen demand fluctuations, supply chain disruptions, and promotional activities. This study contributes to the field of operations management by presenting a novel data-driven approach to characterize inventory policies and ordering behaviors. The findings have practical implications for both retailers seeking to refine their inventory decision-making processes and suppliers aiming to optimize stock planning through more accurate order forecasting.

Enhancing Lasso Selection of Leading Indicators: Predictability and Interpretability

Presenter: Yves R. Sagaert (VIVES University of Applied Sciences)
Author(s): Yves R. Sagaert;Nikolaos Kourentzes

Leading indicators have been shown to be useful in predicting demand. Nowadays, a great number of potentially interesting variables are available in open databases. This makes it challenging to select the predictively relevant indicators. Lasso regression has become a popular choice when handling many variables, however, its selection is sensitive to changes in the sample and in the presence of correlated variables. This can harm both its predictive performance and the trustworthiness of the forecasts. We consider various approaches to aid variable selection. Specifically, we consider clustering via various statistical based methods, semantic information using a Semantic Bidirectional Encoder Representations from Transformers (SBERT), or meta-data such as a popularity index of variables. The resulting groups of variables are evaluated for selection directly, using sequential lasso, or once transformed into cluster profiles, or factors using principal component analysis. Using an inventory management case study, we evaluate the alternative options on their predictive performance, the resulting inventory decisions, and on the interpretability of the resulting models. These are benchmarked against standard lasso and univariate forecasts. Our analysis indicates that there is a trade-off between predictive performance and interpretability.

Forecasting of Inventory Record Inaccuracies

Presenter: Aris Syntetos (Cardiff University)
Author(s): Aris Syntetos;Yacine Rekik;Christoph Glock

A typical retail store contains about 20,000–65,000 distinct stock keeping units (SKUs). To manage such inventories, companies rely upon software tools that use inventory and sales data to automatically predict demand and replenish store shelves. A fundamental assumption of such automated solutions is that our computer inventory records are correct (i.e. what we think we have in stock equals what we actually have in stock). But our research has shown that this assumption is wrong in about 65% of the cases, i.e. more often than not we deal with Inventory Record Inaccuracies (IRI), the economic and environmental implications of which are truly mind blowing. Two types of approaches are currently employed to deal with the problem; and both fail to do so in a cost-effective manner: i) preventive (technology/ RFID-based) and ii) responsive (stock-audit based). However, the wealth of data currently available in retailing enables a new predictive approach, which we have developed with financial support from the Efficient Consumer Response (ECR) and the participation of 9 retailers. Ensemble learning is applied to first predict various classification metrics of IRI, including the binary status of a product (accurate or inaccurate) and the range interval of IRI. These initial classification predictions, achieved with accuracies exceeding 85%, enable the distinction between ‘predictable’ and ‘non-predictable’ sets of SKUs. For the predictable set, a regression-based prediction is then applied to estimate the magnitude of IRI.

Universal Patterns in Cryptocurrency Microstructure: A Machine Learning Approach to High-Frequency Midprice Prediction

Presenter: Bartosz Bieganowski (University of Warsaw)
Author(s): Bartosz Bieganowski;Robert Ślepaczuk

This paper introduces a machine learning model designed to predict the expected value of future midprice changes for high-frequency cryptocurrency trading applications. Our analysis demonstrates remarkable consistency in

feature importance patterns across diverse cryptocurrencies, regardless of market capitalization, trading volume, or tokenomic structure. The model's predictive capabilities enable effective implementation in both taker strategies, which actively execute against existing quotes, and maker strategies, where it signals optimal timing for quote withdrawal. Notably, our findings reveal the universal nature of the learned market relationships, with models achieving peak R^2 performance using as little as one week of training data, even when evaluated against a full year of market activity. These results suggest the presence of fundamental, persistent patterns in cryptocurrency microstructure that transcend specific market conditions and asset characteristics. Our research addresses a critical gap in cryptocurrency trading literature by focusing on short-term price movement predictions between one to three second intervals. Previous work has predominantly emphasized longer time horizons or traditional financial markets, leaving high-frequency cryptocurrency dynamics relatively unexplored. By leveraging a gradient boosting framework trained on order book and trade flow features, our model captures subtle market signals that persist across the cryptocurrency ecosystem. The universality of feature importance rankings across assets as diverse as Bitcoin, Ethereum, and emerging altcoins suggests common underlying mechanisms driving price formation in digital asset markets. This consistency persisted even when testing across bull and bear market cycles, varying volatility regimes, and significant market structure changes like exchange outages or regulatory announcements. Our model maintained predictive power throughout these varied conditions without requiring frequent retraining. Implementation tests showed significant execution cost reductions when integrating model predictions into trading algorithms, with particular effectiveness in high-volatility periods where traditional statistical approaches often fail.

Trading by Charts: A Multivariate CNN System to Predict Retail Investor Trading

Presenter: Chaojie Liu (University of Bristol)
Author(s): Chaojie Liu

Understanding retail investor behaviour is interesting and important but difficult, the prediction of retail investor trading is even much harder. This study sheds some light on predicting retail investors' holding on stocks by building an accessible time-series prediction system using the convolutional neural network (CNN) technique for pre-COVID and COVID-19 periods. A multivariate CNN is built for the first time with both numerical and graphical data. Focusing on the U.S. stock markets between 2018 and 2020, this study utilizes component stocks of the S&P 500 index as the sample with relative data on stock characteristics, retail investor holding, and retail investor ownership. The pioneering multivariate CNN system performs great in predicting retail investor trading and outperforms the random forests models built which only apply numerical data. The results support previous studies on the performances of deep learning techniques like CNN and investor trading behaviour and sentiment. Besides, retail investor holding contains little predictive information for stock price movement. This study contributes to the economic and financial literature by filling the gap in the predictions of retail investor behaviour using cutting-edge machine learning techniques based on novel applications of data. In addition, this prediction system can improve social welfare by helping retail investors make less biased decisions, informing financial institutions to better engage with retail investors, and assisting financial authorities to better monitor and manage risks caused by retail investors in the market.

A Prediction Method for Crude Oil Futures Based on the Optimised Transformer

Presenter: Xiaoxuan Tang (Chengdu University of Technology)
Author(s): Xiaoxuan Tang;Xi Zhang;Yalin Yang

Financial asset price time series typically exhibit characteristics such as high frequency, nonlinearity, volatility, and autocorrelation. To address these challenges, this study proposes a novel crude oil futures price forecasting model based on Variational Mode Decomposition (VMD), Temporal Convolutional Network (TCN), improved Transformer, Bidirectional Long Short-Term Memory Network (BiLSTM), Multilayer Perceptron (MLP), and hybrid architecture—namely VMD-TCN-Transformer-BiLSTM-MLP. In the proposed framework, the Brent crude oil futures closing prices are first decomposed using VMD optimized by Genetic Algorithm (GA). The resulting low-frequency subseries are then fed into the TCN-Transformer-BiLSTM component for modeling, while the high-frequency subseries are input into the MLP. Finally, the predicted values from all subseries are linearly summed to obtain the final forecasting result. The predictive performance of the proposed VMD-TCN-Transformer-BiLSTM-MLP model

is evaluated through error metrics, linear regression tests, MDM tests, and stability analysis. Results indicate that the proposed model significantly outperforms benchmark models such as TCN and BiLSTM in terms of prediction accuracy. Furthermore, precision, recall, and accuracy are employed to assess the model's ability to predict the direction of price movement for the next trading day. The results show that the proposed model also demonstrates strong performance in capturing the trend of Brent crude oil futures prices. Therefore, the model provides valuable insights and decision support for various participants in the crude oil market.

Research on a POS Driven Hybrid Model for Crude Oil Futures Price Prediction

Presenter: Yalin Yang (Chengdu University of Technology)
Author(s): Yalin Ynag;Xi Zhang;Xiaoxuan Tang

As one of the most important commodities in the world, crude oil plays a vital role in the global economic system, and its price fluctuations have far-reaching impacts. Accurate prediction of crude oil prices is crucial for macroeconomic decision-making and risk management. This paper proposes a novel hybrid forecasting model—VMD-BiLSTM-Transformer-MLP—based on PSO to predict WTI crude oil futures prices. First, the PSO algorithm is employed to optimize the parameters of the VMD, which decomposes the crude oil price series into multiple sub-series of different frequencies. The high-frequency sub-series are predicted using a MLP, while the low-frequency sub-series, which exhibit trend characteristics, are processed by a BiLSTM network to capture trend features. A Transformer layer is then used to extract long-term dependencies, followed by another Transformer layer to generate final predictions. The forecasts of all sub-series are linearly combined to obtain the final result. The model's performance is evaluated using RMSE,MAE, MAPE, R^2, and the Modified Diebold-Mariano (MDM) test. Experimental results demonstrate that the proposed model outperforms other benchmark models in forecasting accuracy.

Integrating ENSO Climate Signals into Periodic Autoregressive Models for Wind Speed Forecasting in Brazil

Presenter: Fernando Luiz Cyrino Oliveira (Pontifical Catholic University of Rio de Janeiro)
Author(s): Fernando Luiz Cyrino Oliveira;Rafael Couto;Paula Maçaira

Wind energy has become an increasingly important component of Brazil's renewable energy matrix, particularly in coastal regions. However, existing wind speed forecasting models used by the Brazilian electric sector do not incorporate exogenous climate influences. This study proposes methodological enhancements to the Periodic Autoregressive (PAR) model by including exogenous climate variables—specifically, indicators of the El Niño–Southern Oscillation (ENSO), as well as spatial covariance among states, resulting in PARX and PARX-Cov models. Using reanalysis data from seven Brazilian states and climate forecasts from NOAA and IRI, the models simulate wind speed scenarios across multiple forecast windows. The results demonstrate that incorporating ENSO variables, especially the cumulative Oceanic Niño Index (ONI) and the Southern Oscillation Index (SOI), significantly improves forecast accuracy. In particular, the PARX-Cov model consistently outperformed the benchmark, offering a promising tool for integrating climate signals into renewable energy forecasting and planning in Brazil.

Probabilistic Wind Power Forecasting Based on Multivariable Decomposition and Deep Learning

Presenter: Dongchuan Yang (School of Management, Xi'anXi'an Jiaotong University)
Author(s): Dongchuan Yang;Ju-e Guo;Mingzhu Li

Accurate wind power forecasting is crucial for ensuring the stability and efficiency of power grid operations. However, the inherent volatility of wind power presents significant challenges to precise prediction. This study addresses the limitations of existing decomposition-ensemble forecasting models, which often overlook probabilistic forecasting and underutilize meteorological information. This research introduces a novel decomposition strategy tailored for multivariate time series analysis and designs a specialized deep learning forecasting architecture. This architecture enables the extraction of coupled relationships in both the time and frequency domains from multivariate time series and models the dependencies among decomposed sub-sequences. First, the model employs MVMD to

synchronously decompose wind power and meteorological sequences such as wind speed, enhancing the extraction of inter-variable coupling relationships. An extended dynamic decomposition-reconstruction strategy ensures the effectiveness of the decomposition. Subsequently, the designed deep learning forecasting architecture integrates a multi-head convolutional neural network, a multi-input long short-term memory network with a dual-stage attention mechanism, and a deep ensemble method. These components respectively handle feature enhancement, dynamic weight allocation, temporal correlation capture, and uncertainty quantification. To validate the model's effectiveness, we conducted ablation experiments using wind farm power generation and meteorological data from Yunmeng, Shaanxi Province. We also performed in-depth comparisons of different decomposition strategies and comprehensively analyzed the model's reliability. The experimental results demonstrate that each component of the model contributes to improved forecasting performance. The extended multivariate dynamic decomposition-reconstruction strategy exhibits significant performance advantages compared to six existing multivariate decomposition strategies in the literature. In terms of overall prediction accuracy and reliability, the proposed model outperforms the benchmark models. This study achieves multivariate feature extraction and probabilistic forecasting within a decomposition-ensemble framework, expanding the approaches to wind power forecasting.

Approaches for Probabilistic Wind Power Forecasting

Presenter: Jooyoung Jeon (KAIST (Korea Advanced Institute of Science and Technology))
Author(s): Jooyoung Jeon;Jungyeon Park;Sehyuk Lee

Accurate short-term probabilistic wind power forecasting is crucial for integrating wind energy into power grids while minimizing uncertainty and improving reliability. However, wind power prediction faces significant challenges due to multiple sources of uncertainty, including wind speed variability, the limitations of numerical weather prediction (NWP) models, and the inherent uncertainty in converting wind speed to power generation. To address these challenges, forecasting approaches have traditionally been categorized as direct or indirect methods. This study presents a novel framework that integrates both approaches, incorporating advanced forecasting models and leveraging NWP data and weather variables. The effectiveness of these methods is evaluated across multiple wind farms and compared with conventional approaches.

Spatio-temporal Probabilistic Prediction of Circular Variables: Enhancing the Value of Wind Power Prospective Models by Incorporating Wind Direction Probabilistic Predictions

Presenter: Mario E. Arrieta-Prieto (Universidad Nacional de Colombia)
Author(s): Mario E. Arrieta-Prieto;Kristen R. Schell

Despite the increase in wind farm development over the past decade, the power output of wind farms is still challenging to forecast. In fact, wind power forecast error has increased roughly two-fold over the past five years in the EU grid. Wind direction is known to impact wind power, but the circular nature of this variable has not often been treated correctly in the wind power literature. This work addresses the gap in the literature by developing a model that accounts for its circularity to probabilistically predict wind direction along with wind speed. The impact of this prediction model is then evaluated for its ability to improve the downstream task of probabilistic wind power prediction, for a case study wind farm in Texas. Results show that a GAM + kNN model is the best spatio-temporal model for probabilistic prediction of wind speed and direction. Utilizing the probabilistic predictions of wind speed and direction in the wind power prediction model significantly quantifies the uncertainty in power output. Such a result would be impactful in increasing the amount of wind power dispatched to the grid, in turn helping to lower electricity costs and mitigate climate change.

MTC-MGC: A Multi-Scale Dual-Graph Spatio-Temporal Convolutional Neural Network for Air Quality Prediction

Presenter: Jiaqing Huang (South China Agricultural University)
Author(s): Jiaqing Huang;Dabin Zhang

Accurate air pollution prediction is critical for safeguarding human health, optimizing travel planning, and enhancing environmental management. Despite their transformative potential, existing deep learning-based air quality prediction methods often struggle to effectively model long- and short-term temporal dependencies and dynamic spatial correlations. To address these challenges, this study introduces a novel air quality prediction model, MTC-MGC, which integrates a Multi-scale Temporal Convolution (MTC) module and a Mix-hop Graph Convolution (MGC) module. The MTC module employs multi-kernel dilation convolution to efficiently capture both local and global temporal patterns in air quality data. By adaptively learning hierarchical temporal features at multiple time scales, this mechanism significantly enhances the model's ability to model complex temporal dependencies inherent in air quality datasets. The module leverages convolutional kernels of varying sizes to simultaneously capture fine-grained short-term fluctuations and broader long-term trends, addressing the common issue of temporal information loss in traditional convolutional architectures. To further improve spatial modeling, the MTC-MGC incorporates a Mix-hop Graph Convolution module. This module extracts stable long-term spatial correlations using an adaptive static graph, while generating dynamic frequency-domain graph based on Fourier-transform-derived frequency information and raw spatio-temporal data. This dual-graph structure effectively captures both stable and evolving spatial relationships among monitoring sites, This dual-graph structure can effectively extract stable as well as changing spatial relationships between sites, which can present the spatial dependencies in the air quality data more comprehensively and improve the ability of the model to simulate the air pollutant dispersion process. Extensive experiments conducted on three real-world air pollution datasets—Guangdong Province, Yangtze River Delta region, and Beijing-Tianjin-Hebei region—demonstrate that the proposed MTC-MGC model outperforms state-of-the-art baselines in terms of prediction accuracy and trend fitting. These results validate the model's effectiveness and robustness in handling complex environmental data. This study offers a robust technical framework for air quality prediction, providing strong support for urban environmental management and public health decision-making.

Carbon Price Forecasting Based on Multi-level Features and Bayesian Optimized LSTM

Presenter: Dong Zeng (South China Agricultural University)
Author(s): Dong Zeng;Huanling Hu;Dabin Zhang

The carbon trading market is a core policy tool for achieving the "dual carbon" goal. Its price fluctuations have significant nonlinear and multiple uncertainties, which brings great challenges to accurate prediction. This study proposes a carbon price prediction method that integrates multi-feature enhancement and optimization deep learning models, targeting the nonlinear and dynamic characteristics of carbon price data. First, the sliding window technology is used to segment the historical closing prices and construct the initial feature set. Then, the probability density network (PDN) and random walk similarity (SRW) bimodal analysis methods are innovatively combined: PDN quantifies the distribution similarity between time windows, and constructs a weighted network to characterize the implicit pattern of the carbon price series. SRW is used to measure the dynamic evolution similarity of the time series, characterize the complex dependencies of the data, and the two heterogeneous features are fused to construct an enhanced feature matrix. On this basis, a Bayesian optimized LSTM network is designed, and the intelligent optimization of network hyperparameters is realized through the Gaussian process surrogate model, which effectively solves the problem of traditional LSTM parameter settings relying on experience. Experimental results show that the prediction accuracy of this method is significantly improved compared with the benchmark model on the EU ETS dataset, exceeding the effectiveness of other methods.

A Spatiotemporal Forecasting Model Based on Local and Global Spatial Correlations for Multi-node Wind Speed

Presenter: Ni Hong (South China Agricultural University)
Author(s): Ni Hong;Dabin Zhang

Accurate wind speed forecasting is crucial for optimizing wind energy utilization and scheduling wind power generation. However, most existing methods rely on prior knowledge to construct spatiotemporal graph structures for feature extraction, or jointly model spatial and temporal dependencies through separate spatial and temporal modules, limiting their ability to learn intricate and dynamic spatiotemporal patterns directly from data. To address this issue, we propose a novel spatiotemporal prediction model, namely Local and Global spatial Graphs combined for Wind

speed forecasting Network (LGG-WindNet), which captures the dynamic wind speed patterns across multiple nodes by jointly modeling local and global spatial correlations. First, the number of neighboring nodes in the graph is constrained using Dynamic Time Warping (DTW) to enforce sparsity, and then a local convolutional unit is employed to capture strong local dependencies among wind speed nodes. Subsequently, a global self-attention module is applied to model long-range dependencies across distant nodes in the wind field. Additionally, graph regularization is utilized to control the quality of the learned graph structure. Second, the model incorporates dynamic temporal features of nodes to models long-term temporal dependencies, thereby accounting for the time-varying characteristics of relationships between different time points due to changing weather conditions. Third, within an end-to-end framework, the proposed model integrates local and global graph learning, temporal convolutional networks (TCN), and graph convolutional networks (GCN) to jointly model the nonlinear temporal dependencies and dynamic spatial structures of wind speed. Finally, experiments conducted on two real-world wind farm datasets demonstrate that LGG-WindNet significantly improves wind speed forecasting accuracy by effectively leveraging spatiotemporal features.

Calibrated Multi-Level Quantile Forecasting

Presenter: Tiffany Ding (University of California, Berkeley)
Author(s): Tiffany Ding;Isaac Gibbs;Ryan Tibshirani

In order for probabilistic forecasts to be useful to decision makers, the forecasts should be calibrated – given a sequence of 90% quantile forecasts, we want the true value to be less than the forecast 90% of the time. One way to calibrate a sequential forecaster is to use the quantile tracking (QT) algorithm (Angelopoulos et al., 2023) from online conformal prediction. This algorithm is guaranteed to achieve calibration in the long run. However, in many settings, we want to generate forecasts for multiple quantiles simultaneously. Unfortunately, applying QT to each quantile often leads to crossing – e.g., the calibrated 50% quantile forecast is above the calibrated 75% quantile forecast. This is very undesirable from the perspective of practitioners. This leads to our motivation question: how can we produce quantile forecasts for multiple quantile levels that achieve the correct coverage while not crossing? We answer this problem from both theoretical and empirical perspectives. We propose intuitive ways of combining the quantile tracking algorithm with an order-enforcing method (such as sorting or isotonic regression) that produces a sequence of forecasts with no crossings but is also guaranteed to achieve the correct long-run coverage under mild assumptions. We demonstrate our methods on COVID-19 forecasting data.

Conformal Prediction Under Model Uncertainty

Presenter: Xiaodong Yan (西安交通大学)
Author(s): Xiaodong Yan

Model averaging is a robust approach for dealing with model uncertainty, however, constructing prediction intervals through model averaging can be challenging due to the allocation of weights. To address this challenge, we propose a novel approach called the conformalized jackknife model averaging (CJMA) method, which combines model averaging and conformal prediction and utilizes the jackknife criterion. Within this framework, each candidate model contributes to the conformal prediction process, and weights calculated from residuals are allocated to each model for aggregation. By leveraging the strengths of both model averaging and conformal prediction, the CJMA method enables the generation of efficient prediction intervals. It has the advantage of producing shorter confidence intervals while maintaining the same target coverage. Additionally, the CJMA method outperforms the jackknife conformal prediction when handling out-of-sample predictions. We demonstrate the performance of the CJMA method through extensive simulation studies and its application to a rat eye microarray expression dataset.

Empirical Study on the Performance of Conformal Forecasting Methods on Real-world Data

Presenter: Roberto Morales-Arsenal (CUNEF Universidad)
Author(s): Roberto Morales-Arsenal;Christoph Bergmeir;Lorenzo Escot-Mangas

Conformal prediction is a statistical framework that generates valid predictions with quantifiable uncertainty,

ensuring reliability in fields like machine learning. By leveraging exchangeability, it remains statistically valid even in finite samples, producing prediction intervals under minimal assumptions, which is crucial for high-stakes decisions. Despite its advantages, conformal prediction faces several challenges, including the need for large calibration datasets and difficulties in scenarios with non-exchangeable data, such as time series forecasting or dynamically evolving datasets. In such cases, traditional conformal prediction methods require adaptations to handle temporal dependencies.This study investigates the impact of violating these assumptions on coverage and the width of conformal intervals, extending the analysis to a diverse set of time series with different characteristics, such as varying levels of autocorrelation, non-stationarity, and structural changes.

Probabilistic Bake-off: How to Quantify Uncertainty with Neural Networks

Presenter: Olivier Sprangers (Nixtla)
Author(s): Olivier Sprangers

Probabilistic forecasts are vital in many application domains where practitioners are in need of a notion of the uncertainty of the forecast. The surge in neural network architectures for forecasting has also sparked interest in their ability to create probabilistic forecasts. Nowadays, many methods are available to create such probabilistic forecasts. For example, one can use (i) a parametric distribution where the parameters of a fixed distribution are learned by the neural network, or (ii) a (weighted) mixture distribution where the network learns to use a set of parametrized distributions jointly, or (iii) a set of quantile outputs where the network learns specified output quantiles directly, or (iv) conformal prediction intervals, and many more. We believe that a rigorous benchmark across methods, architectures and datasets is currently lacking in the forecasting field. In this talk, we benchmark over 20 different neural network architectures comprising RNN, CNN and Transformer-based univariate and multivariate forecasting architectures on their ability to produce probabilistic forecasts using contemporary probabilistic output methods. In addition, we validate against machine-learning and statistical benchmark methods. An interesting result is that conformal methods – despite their theoretical benefits – do not necessarily offer the best probabilistic forecasting results in practice. After this talk, a forecasting practitioner has an understanding of which methods are available to quantify uncertainty when forecasting with neural networks and which methods perform well in which scenarios.

Factors Influencing Medical Resource Congestion during Infectious Disease Outbreaks: Analysis Based on a Medical Resource Dynamics Model

Presenter: Jiaming Guo (Shanxi Medical University)
Author(s): Jiaming Guo

Infectious disease outbreaks challenge medical resource allocation as the resultant surges in hospitalizations and medical technician and bed shortages prevent patients from receiving urgent care, which further exacerbates public health crises. Using COVID-19 data from Taiyuan and Jincheng (from December 9, 2022–January 13, 2023), we explored the impact of medical resource allocation on the onset and duration of medical resource congestion to provide a reference for future epidemic prevention and control measures. We measured medical resource congestion using the number of available general medical technicians, intensive care unit (ICU) beds, and ICU medical technicians. Based on the dynamic relationships between hospital demand and medical resource supply, we then established a dynamical model to analyze the factors influencing resource congestion during COVID-19. The results revealed that the medical resource dynamics model had good fit; medical resource congestion occurred during the pandemic; the maximum values of the resource allocation, allocation rate, and onset time of the resource allocation impacted the congestion onset and duration; and Taiyuan and Jincheng differed in terms of the congestion onset and duration. During future epidemics, it is necessary to combine different regional characteristics and focus on the allocation of general medical technicians and allocation rates of ICU beds and ICU medical technicians to alleviate medical resource congestion duration.

MuSTANet: A Multi-Scale Trend-Aware Network with Dynamic Variable Selection and Squeeze-Excitation for Interpretable Carbon Allowance Returns Forecasting

Presenter: Dinggao Liu (Fujian Agriculture and Forestry University)
Author(s): Dinggao Liu;Zhenpeng Tang;Yi Cai

Accurate carbon allowance returns forecasting is crucial for policymakers, investors, and industries seeking to effectively manage carbon emissions and hedge against market uncertainties. However, predicting carbon allowance returns faces the challenges of volatile short-term fluctuations, longer-term trend shifts, and the heterogeneous influence of multiple related market factors. Meanwhile, the traditional multi-head self-attention mechanism was initially proposed to process the discrete tokens and fails to consider the local trend information inherent in continuous data. In this paper, we propose a multi-scale trend-aware network (MuSTANet) to capture and match temporal patterns in complex carbon allowance returns sequences. Specifically, we introduce a multi-scale trend-aware self-attention mechanism that leverages parallel convolutional branches with varied receptive fields to discern local slope changes and match temporally relevant points. Furthermore, a dynamic variable selection strategy that adaptively reweights input features at each time step, enabling the model to focus on salient indicators—such as energy prices and macroeconomic signals—when their impact on returns is most pronounced. Multi-scale convolutions then extract temporal patterns at varying granularities to capture both rapid fluctuations and persistent trends, with the squeeze-and-excitation block recalibrating channel-wise responses to amplify important feature channels while suppressing less informative ones. We conduct extensive experiments on 4 carbon trading systems, including the European Union Emissions Trading System (EU ETS) and different Chinese carbon markets. The results demonstrate that MuSTANet improves predictive accuracy, directional accuracy, and achieves the best performance in trading backtests against previous competitive methods. Moreover, MuSTANet provides interpretable insights into the role of each variable at different time horizons. Its adaptability and interpretability not only provide robust forecasting tools for policymakers and investors but also furnish a more scientific basis for risk management and investment decisions in carbon markets.

Can Online Environment Public Opinion Promote the Enterprises' green Technology Innovation? –An Empirical Research Based on Baidu Index

Presenter: Yuxin Ran (Beijing Technology and Business University)
Author(s): Yuxin Ran;Yanmin Shao;huanqi Shi

The rapid development of the Internet has expanded the public's access to information and expression, and has provided them with opportunities to concern increasingly serious environmental issues. However, the impact of online environment public opinion(OEPO) on enterprise green technology innovation (GTI) has been under-explored. Drawing on organizational legitimacy theory, we posit that an increased OEPO exerts organizational legitimacy pressure on enterprises, driving their participation in GTI. Further, we argue that pollutant emission information disclosure and supportive policy attitude will strengthen the relationship by amplifying organizational legitimacy pressure. Using the regional Baidu index for environment-related online public keywords to capture OEPO, we obtain empirical evidence supporting our arguments from a sample of Chinese listed enterprises. As the first study exploring the relationship between OEPO and enterprise GTI, this study may inspire regulators to leverage the power of the public to promote enterprise green development and achieve high-quality development of the national economy.

A Novel Disaster Index to Support Pandemic Planning: A Principal Component Analysis-Based Approach

Presenter: Alessia Paccagnini (University College Dublin)
Author(s): Alessia Paccagnini;Annunziata Esposito Amideo;Luca Pistilli

COVID-19 has severely disrupted everyday life aspects of communities around the world. Being at stake health, society, and economics worldwide, we propose an interdisciplinary approach to support pandemic planning. Specifically, we integrate economics , management and business analytics, to a greater extent and apply such disciplines to disaster management We present two novel disaster indexes capturing various facets such as economic

GDP and health-related factors. In particular, based on the Disaster Operations Management framework, we define a pre-disaster evaluation index (Pre-DMEi) and a post-disaster evaluation index (Post-DMEi). The former provides the snapshot of a country's economic and health status before the occurrence of a disaster while the latter captures the reaction to the disaster occurrence thus offering a dynamic evaluation of the disaster response. The pairing of these two indexes allows one to gain insight on a country's pre-disaster status and its response. Moreover, Post-DMEi allows identifying what aspects to invest in to mitigate adverse effects if such a disaster shall occur again: this is achieved through Principal Component Analysis (PCA). For this study, we consider COVID-19 as the disaster and we test our framework on countries within Europe. e then conduct an out-of-sample forecasting analysis, assessing the predictive power of the Pre-DMEi in forecasting post-pandemic economic and healthcare recovery trends. This forecasting exercise allows us to test the stability of our framework and its applicability in future crises. This research aims to help in the evaluation, measurement, and improvement of the set of interventions needed to withstand COVID-19 and, most importantly, aid policy-makers in light of new pandemics in the time to come.

Research on the Path of New Quality Productivity Driving Household Indirect Carbon Emission Reduction under the "Dual Carbon" Goal

Presenter: Zirui Chen (Xidian University)
Author(s): Zirui Chen

In the context of global response to climate change and China's "dual carbon" goals. As an important source of household consumption, household indirect carbon emission and its emission reduction path needs to be explored urgently. As the core driving force for green transformation, It is of great significance to discuss the energy consumption and carbon emission status of the household, an important energy consumption sector. Based on the China Family Panel Studies (CFPS) in 2016, 2018 and 2020, this paper constructs a two-way fixed-effect model to empirically analyze the impact of new quality productivity on indirect carbon emissions in households; on this basis, the mediation analysis model is used to verify the mechanism of the new quality productivity on indirect carbon emissions in households; and further explores the heterogeneity of new quality productivity on indirect carbon emissions in households based on the perspective of family regions and income levels. The results show that: First, the improvement of new quality productivity significantly reduces the indirect carbon emission level of households; Second, the level of Internet development is an important path for new quality productivity to affect household carbon emissions; Third, the emission reduction effect of new quality productivity is more obvious in the central and western regions and high-income families, mainly due to regional resource optimization and the advantages of green technology application for high-income groups. Research suggests that policy design needs to combine regional resource endowments and residents' income differences, strengthen the synergy between new quality productivity and digital technology, and promote low-carbon transformation of families.

Systemic Risks and Contagion Effects in China's Regional Electricity Exchange Network: A Differential DebtRank Model

Presenter: Yue Pan (Xidian University)
Author(s): Yue Pan

Ensuring the safe and stable operation of the electricity system has long been a priority in building China's new electricity system and is crucial for driving energy transformation and supporting high-quality economic and social development. However, few studies have examined the systemic risks arising from the complex electricity transmission relationships between different regions in China. To address this, we developed a differential DebtRank model from the perspective of regional electricity interconnection, considering the existence of backup coal-fired electricity units, to simulate and analyze the systemic risks and contagion effects posed by each region as a risk source. Additionally, to better understand the significant differences in systemic risks across regions, we used panel data from 30 provinces and municipalities in China (2016-2022) to examine the factors and mechanisms influencing systemic risks in regional electricity exchanges. We found that China's electricity exchange network has a certain resilience to risks, with strong interdependence and complementarity among its nodes. In 2022, Inner Mongolia and Ningxia emerged as key nodes, potentially causing the greatest systemic risks. Conversely, economically developed regions like Beijing and Shanghai have the lowest systemic risks. According to the six regional grids, risk contagion mainly

occurs within individual grids, with the northwest power grid posing the greatest national impact. Using the "Hu Line" as a boundary, systemic risk is higher in the northwest, lower in the southeast, and the gap is widening over time. Empirical analysis reveals that regions with high GDP, strict environmental regulations, and strong transportation infrastructure face lower systemic risks, whereas electricity generation scale and net export value are key factors in amplifying these risks. Our research not only offers a new analytical framework for electricity sector risks but also provides theoretical support and decision-making references to enhance the security resilience of China's electricity network and promote energy transformation

How Do Smart Aging-Friendly Policies Affect Elderly Households' Energy Consumption? Causal Evidence and Policy Scenario Simulations from Intergenerational Technology

Presenter: Jiang Yichun (Xidian University)
Author(s): Jiang Yichun

Rapid population aging poses significant challenges for China, balancing improved elderly welfare with residential carbon reduction. This paper empirically investigates how smart aging-friendly policies influence elderly households' energy consumption by facilitating intergenerational digital technology transfer. Using China's nationwide aging-friendly subsidy program as a quasi-natural experiment, we develop a spatial-temporal double machine learning framework to analyze integrated datasets based on CHARLS. Results suggest that aging-friendly subsidies effectively reduce household energy intensity, an effect notably strengthened by intergenerational digital mentoring and supportive community digital infrastructure. Policy scenario modeling through system dynamics further indicates that carefully designed policy packages—combining targeted subsidies with local digital infrastructure upgrades—can significantly reduce household emissions. However, increased comfort-driven energy use (rebound effects) may partially offset gains, underscoring the need for complementary mitigation strategies. Our study provides robust causal evidence on the key role intergenerational relationships play in household energy transitions, alongside practical insights for policymakers pursuing sustainable aging policies.

Enhancing Value-at-Risk Forecasting: A Hybrid Deep Learning Approach Leveraging Textual Data

Presenter: Yangfan Cao (sunway univeristy)
Author(s): Yangfan Cao;Xi Luo

The growing adoption of deep learning and natural language processing (NLP) has significantly enhanced efficiency across various domains. This study aims to improve the accuracy of Value-at-Risk (VaR) models forecasting by integrating text mining and deep learning techniques. It employs NLP to analyze the sentiment and subjectivity of online news headlines, incorporating them as variables to refine stock market risk forecasts and assess their influence on VaR accuracy. Additionally, the study utilizes the Latent Dirichlet Allocation (LDA) topic model to uncover hidden topics in news articles, enriching the textual information used in VaR model construction. To further enhance predictive performance, it integrates forecasts from four Generalized AutoRegressive Conditional Heteroskedasticity (GARCH)-type models with advanced Long Short-Term Memory (LSTM) and Convolutional Neural Network (CNN)-LSTM models. The study also examines the impact of textual data on VaR predictions over different timeframes, specifically using 7-day and 20-day rolling windows. Using historical data from the S&P 500 (SPY), Dow Jones Industrial Average (DJI), and Nasdaq Composite (IXIC) from 2012 to 2023, along with corresponding news headlines, this research rigorously evaluates the proposed methodology. The findings demonstrate that incorporating textual data into VaR models significantly enhances forecasting accuracy, highlighting the advantages of deep learning in financial risk prediction. Moreover, longer rolling windows improve VaR estimation, likely due to CNN's ability to extract valuable textual insights and LSTM's effectiveness in capturing both short-term and long-term dependencies.

Forecasting and Monitoring Cryptocurrency Risk: Evidence from a Composite Risk Index

Presenter: XC Guo (University of Science and Technology Beijing)
Author(s): Xiaochun Guo;Shouyang Wang;Kun Guo

This paper proposes a new framework for forecasting and monitoring risk in cryptocurrency markets through the development of the Crypto Risk Composite Index (CRCI). The CRCI captures four key dimensions of systemic risk: market behavior, network activity, investor sentiment, and relative performance against traditional assets. An entropy-CRITIC weighting method is employed to objectively integrate heterogeneous data sources, offering a dynamic and data-driven view of evolving risk profiles. We apply CRCI to a dataset covering January 2015 to February 2024, during which the index successfully identifies four major high-risk episodes: the late-2017 ICO bubble burst, the early 2020 COVID-19 shock, the 2021 speculative cycle, and the 2022 correction. Toward the end of the sample, CRCI signals a renewed upward trend in systemic risk, suggesting increasing market fragility. To evaluate the index's forecasting capacity, we compare CRCI with the Value-at-Risk (VaR) of major crypto assets and wildly used crypto indices. We assess their correlations with decomposed risk components via Dynamic Conditional Correlation (DCC) models. The results demonstrate that CRCI not only anticipates spikes in VaR but also captures shifts in dominant risk drivers. The result shows that traditional asset comparisons became more influential after COVID-19 outbreak, while sentiment and uncertainty exhibited a growing role. Network activity, though generally stabilizing, occasionally served as a leading indicator of risk escalation. Unlike existing crypto indices, which primarily reflect market capitalization and structural changes, CRCI is explicitly designed for risk evaluation and forecasting. Its multidimensional framework enables both aggregate and source-specific risk tracking. By combining entropy and CRITIC in the weighting process, CRCI innovatively integrates diverse risk factors into a unified index. Overall, CRCI offers a forward-looking, granular perspective on cryptocurrency market risk with strong potential for early warning applications in investment strategy, portfolio management, and regulatory oversight.

On the Implied Equity Risk Premium

Presenter: Shuyuan Qi (Central University of Finance and Economics)
Author(s): Shuyuan Qi;Emese Lazar;Radu Tunaru

This study advances the literature on forward-looking equity risk premia through two pivotal contributions. We focus on the implied equity risk premium, as a function of investors' risk aversion and risk-neutral return moments, circumventing the inherent difficulties associated with estimating higher-order physical return moments, thereby establishing a more robust analytical framework for the computation of forward-looking equity risk premia. We demonstrate that the skewness spread, in addition to the volatility spread, exerts a significant influence on investors' risk aversion and on the implied equity risk premia term structure. We augment the existing risk aversion estimation paradigm by integrating both volatility and skewness spreads as orthogonality conditions within a generalized method of moments estimation framework. Our examination of the implied risk premia term structure reveals different behaviours before, during, and after market crises. Furthermore, the new equity premium measure shows great performance in asset allocation exercises and it demonstrates remarkable predictive capacity for future real economic activities across varying forecasting horizons.

Forecast for Regional Air Flow

Presenter: Jiancong Huang (Middle South Air Traffic Management Bureau)
Author(s): Jiancong Huang

At present, the prediction of regional airspace flow is mostly based on the average value of accumulated historical data, and is evaluated using only a single dimension, lacking consideration for dynamic adjustments and macro level factors. However, the error is very giant. The forecast results can only be published near the final deadline. This article considers the changes in airport capacity, sector capacity, newly opened routes, airspace, department and arrival procedures, and other elements in different years compared to the previous period, as well as the impact factors of seasonal changes on different regions at the meteorological level. At the same time, the importance of traffic flow varies among the main coordination airports, auxiliary coordination airports, and other transportation airports. It is necessary to calculate weights separately and estimate the impact of flight on-time performance rates on execution rates based on the number of flight plans at each airport. In terms of human factors, it is essential to consider the impact of other airspace user activities in the past few years and calculate the impact weights based on the flow control of different months. In other aspects, alternative modes of transportation, such as new high-speed rail lines, may replace the past way of transportation and require further analysis. According to the time span, the weights of former

factors are different among one month, one quarter, and the whole year. This data can be referred as the arrangement of duty personnel, dispatch of aircraft stands, release of flow control, and allocation of airspace capacity. To some extent, through the relevant forecast of the transportation airports, the forecast analysis method mentioned in this article can be used for forecasting the flow of low altitude flights such as general aviation and drones in the future.

Transformer-Based Emotion Intensity Analysis for High-Frequency Prediction of Market Overreactions and Trading Strategies

Presenter: Szymon Lis (University of Warsaw)
Author(s): Szymon Lis;Robert Ślepaczuk;Paweł Sakowski

This study investigates the application of transformer-based emotion intensity analysis in high-frequency trading (HFT) to detect and exploit short-term market overreactions. We leverage the DistilRoBERTa model to classify real-time Twitter sentiment related to Apple Inc. (AAPL) and construct an emotion index reflecting shifts in investor sentiment. By integrating this index into an HFT framework, we generate predictive trading signals that capitalize on sentiment-driven mispricing, either by following momentum-based trends or exploiting contrarian mean reversions. We develop and test three sentiment-driven trading strategies: (1) a sentiment momentum strategy, which assumes that strong emotion-driven price movements exhibit short-term continuation; (2) an emotion-contrarian strategy, which seeks to exploit extreme investor overreactions that revert after sharp sentiment-induced price swings; and (3) a high-intensity filtered strategy, which incorporates engagement metrics (e.g., retweet volume, likes) to refine signal reliability and reduce noise. Backtesting was conducted on high-frequency AAPL market data from January to June 2022. The high-intensity filtered strategy outperformed both other sentiment-based approaches and a volume-weighted average price (VWAP) benchmark, achieving the highest Sharpe ratio (1.68) and win rate (56.1%). The results suggest that filtering sentiment signals based on intensity and engagement improves predictive accuracy, enhancing trade execution and profitability. To optimize model parameters, we employ machine learning techniques, including LSTM, to calibrate key trading parameters such as sentiment thresholds, trade execution timing, and volatility-adjusted risk controls. Adaptive deep learning models are used to dynamically adjust sentiment-driven entry and exit signals, further refining trade precision. Our findings highlight the effectiveness of transformer-based sentiment analysis in financial markets, demonstrating its potential to enhance algorithmic trading strategies. However, results also underscore the necessity of risk management techniques, including trade frequency constraints and volatility-sensitive execution filters, to mitigate overtrading and minimize exposure to transient sentiment shocks. Our study contributes to the growing field of AI-driven trading by demonstrating the value of real-time emotion analysis in high-frequency financial decision-making. Future

Context-Driven Cold-Start Web Traffic Forecasting

Presenter: Xin Zhou (Monash University)
Author(s): Xin Zhou;Weiqing Wang;Wray Buntine;Christoph Bergmeir

Cold-start forecasting is critical in dynamic scenarios where early-stage forecasting drives key decisions, such as content prioritization, resource allocation, and demand estimation before observable trends emerge. In this work, we explore the potential of multimodal forecasting techniques for cold-start forecasting and offer insights into designing more scalable and adaptive models. In particular, we address context-driven cold-start web traffic forecasting that includes textual content and historical web traffic of relevant web pages to generate forecasts when no historical data are available for the target new web page. To advance research in this area, we collect, clean, and align a high-dimensional, multimodal web traffic dataset. We adopt a Retrieval-Augmented Generation framework, and propose the use of large language models (LLMs) for this task. Our experiments demonstrate that the LLM-based strategy consistently outperforms the statistical baseline across multiple forecasting horizons. The best-performing LLM-based model reduces WRMSPE by 10.8% and WAPE by 18.2%, compared with the statistical baseline. Furthermore, LLM-based feature extraction enhances contextual understanding, leading to greater stability in long-horizon forecasts.

Coherent Forecast Combination: a Stacked Regression Approach

Presenter: Daniele Girolimetto (University of Padova)
Author(s): Daniele Girolimetto;Tommaso Di Fonzo

Forecast reconciliation is a well-established post-forecasting process that adjusts base forecasts produced by a single expert or model to ensure coherence within a constrained forecasting framework. However, in many practical applications, multiple forecasts are available for each variable, originating from different models or experts. Coherent forecast combination extends traditional reconciliation by integrating forecast combination and reconciliation into a unified framework to improve accuracy and satisfy the coherence property. This methodology includes both optimal (in the least squares sense) and sequential approaches. In the former case, forecast combination and reconciliation occur simultaneously, while in the latter either step is performed separately: first single-variable forecast combination is performed, and then the resulting forecasts are reconciled, or vice versa. In this talk, we explore the theoretical principals of coherent forecast combination, its advantages over single-task combination and single-expert reconciliation approaches, and its practical implementation with the R package FoCo2. Using a real-world energy dataset, we illustrate its effectiveness and its ability to improve forecasting accuracy while fulfilling constraints.

Score-optimal Forecast Reconciliation

Presenter: Shun Hu (Beihang University)
Author(s): Shun Hu;Bohan Zhang;Yanfei Kang

Hierarchical time series are commonly encountered in various domains, prompting significant advancements in forecast reconciliation techniques over the past decade. While the well-established Minimum Trace method has been effective in enhancing forecast accuracy by considering covariance information, it fails to account for other critical distribution information. Real-world hierarchical time series often exhibit varied characteristics in different levels, resulting in mixed-type forecast distributions that can be normal, non-normal, or non-parametric. In response to these limitations, this paper introduces a novel forecast reconciliation method designed to optimize the coherent point forecast based on a user-specified scoring rule, thereby capturing comprehensive information from the given forecast distributions. This method incorporates a Gaussian-copula based sampling technique, making it applicable to any type of continuous distribution. This approach further produces coherent probabilistic forecasts by employing a recalibration process. We demonstrate the robustness of our method through a three-node example featuring non-normal distributions, highlighting its superiority over the Minimum Trace approach. Additionally, the proposed approach is applied to day-ahead half-hourly forecasts for load, wind, and solar power generation in Shanxi province of China, showcasing its practical effectiveness and adaptability.

A Network Regression Approach for Forecast Reconciliation

Presenter: Jiaxin Shi (Peking University)
Author(s): Jiaxin Shi;Feng Li;Hansheng Wang

Forecast reconciliation plays a crucial role in hierarchical and grouped time series forecasting, ensuring coherence across different levels of aggregation. While existing reconciliation methods have shown promising empirical performance, their statistical properties remain less explored. In this paper, we propose a network regression approach for forecast reconciliation, which provides a flexible framework for modeling dependencies among forecasts. Specifically, we utilize covariance regression methods to estimate the covariance structure of forecast errors. The unknown parameters are estimated by the quasi-maximum likelihood estimation method, and we establish the theoretical properties of the proposed method, showing its desirable statistical performance. Empirical studies on both simulated and real-world datasets demonstrate the finite sample performance of the proposed method.

Graph Signal Processing for Global Stock Market Realized Volatility Forecasting

Presenter: Zhengyang Chi (University of Sydney)
Author(s): Zhengyang Chi;Junbin Gao;Chao Wang

The interconnectedness of global financial markets has brought increasing attention to modeling volatility spillover effects in volatility forecasting tasks. This paper introduces an innovative realized volatility (RV) forecasting framework that extends the conventional Heterogeneous Auto-Regressive (HAR) model via integrating the Graph Signal Processing (GSP) technique. The proposed framework can be referred to as the GSPHAR framework. The volatility spillover effect is embedded and modeled in the GSPHAR framework, which employs the graph Fourier transformation method to effectively analyze the global stock market dynamics in the spectral domain. In addition, convolution filters with learnable weights are applied to capture the historical mid-term and long-term volatility patterns. Furthermore, the proposed framework is capable of incorporating dynamic graph structures into the proposed RV forecasting model (the GSPHAR model) to capture the changing volatility interrelationships. The empirical study is conducted with RV data of 24 global stock market indices with around 3500 common trading days from May 2002 to June 2022. The short-term, middle-term and long-term RV forecasting performance of the GSPHAR model is compared with various HAR type models. The results show that the proposed model consistently outperforms all other models considered in the study, demonstrating the effectiveness of integrating the GSP technique into the HAR model for RV forecasting.

Adaptive Online Bagging with Hybrid Sampling for Real-time Credit Card Fraud Detection

Presenter: Yiming Teng (Harbin Engineering University)
Author(s): Yiming Teng;Lean Yu;Hang Yin

Handling imbalanced data streams with evolving distributions remains a critical challenge in real-time credit card fraud detection, where concept drift and class imbalance degrade model performance over time. An Adaptive Online Bagging with Hybrid Sampling (AOBHS) approach is proposed to address this issue. AOBHS combines active drift detection and passive drift adaptation to manage evolving data distributions, while mitigating class imbalance through a hybrid sampling technique. This technique includes enhanced online bagging extensions that respond to past drift events and leverage rare samples, along with a novel SMOTE-variant oversampling method optimized using sketch distribution. Experiments on synthetic data streams and real-world credit card fraud datasets show that AOBHS outperforms benchmark methods in terms of Prequential AUC, G-mean, and recall, demonstrating its effectiveness for real-time fraud detection.

Transfer Learning Under High-Dimensional Network Convolutional Regression Model

Presenter: Danyang Huang (Renmin University of China)
Author(s): Danyang Huang

Transfer learning enhances model performance by utilizing knowledge from related domains, particularly when labeled data is scarce. While existing research addresses transfer learning under various distribution shifts in independent settings, handling dependencies in networked data remains challenging. To address this challenge, we propose a high-dimensional transfer learning framework based on network convolutional regression (NCR), inspired by the success of graph convolutional networks (GCNs). The NCR model incorporates random network structure by allowing each node's response to depend on its features and the aggregated features of its neighbors, capturing local dependencies effectively. Our methodology includes a two-step transfer learning algorithm that addresses domain shift between source and target networks, along with a source detection mechanism to identify informative domains. Theoretically, we analyze the lasso estimator in the context of a random graph based on the Erdős–Rényi model assumption, demonstrating that transfer learning improves convergence rates when informative sources are present. Empirical evaluations, including simulations and a real-world application using Sina Weibo data, demonstrate substantial improvements in prediction accuracy, particularly when labeled data in the target domain is limited.

Energy Independence and Economic Resilience

Presenter: Jiayin Sun (Renmin University of China)
Author(s): Jiayin Sun

This study examines how energy independence influences a country's economic resilience, stability, and capac-

ity for sustainable development, in response to the growing challenges posed by global geopolitical tensions, energy supply uncertainty, and volatile energy prices.Empirical results indicate that energy-independent countries exhibit significantly greater resilience to international shocks. Specifically, the contribution of global shocks to both output and inflation volatility is markedly lower in energy-independent economies compared to energy-dependent ones.Furthermore, the transformation of national energy structures and improvements in energy self-sufficiency are shown to play a crucial role in mitigating external shocks, stabilizing macroeconomic performance, and enhancing the effectiveness of policy design and risk management.

Perceived Uncertainty and the Effects of Monetary Policy in China

Presenter: Ningxin Qiu (School of Applied Economics, Renmin University of China)
Author(s): Ningxin Qiu;Haibo Li;Xinye Zheng;Zidong An

As a higher-order state of economic expectations, perceived uncertainty profoundly impacts the macroeconomy. In the current context of weak social expectations and increasing uncertainty in China, it is of great significance to grasp the perceived uncertainty and its policy implications to strengthen the effects of macroeconomic policy. Based on micro-level economic forecast data, this paper measures perceived uncertainty from 1996 to 2022 in China and empirically examines its impact on the effects of monetary policy. Compared with macroeconomic uncertainty, which reflects the overall shocks to the economy, and economic policy uncertainty, which reflects the policy-relevant part of the shocks, perceived uncertainty directly reflects economic agents' responses. This paper finds that the perceived uncertainty in China rises when the economy faces a shock, significantly weakening the effects of monetary policy. These findings emphasize the role of expectation in monetary policy transmission and provide an alternative explanation of the heterogeneous monetary policy effectiveness across different periods. Under the constraint of limited policy space, the central bank can enhance the effect of monetary policy by reducing the uncertainty of economic expectations.

An Instrumental Variables Approach to Testing Forecast Efficiency

Presenter: Xuguang Simon Sheng (American University)
Author(s): Xuguang Simon Sheng;Tucker McElroy

To achieve forecast efficiency, it is necessary that forecast errors are unpredictable from forecast revisions. Coibion and Gorodnichenko (2015, AER) propose aggregating forecasts and estimating the aggregated time series regression. Bordalo, Gennaioli, Ma, and Shleifer (2020, AER) suggest estimating the average relationship by running separate regressions for each individual and then aggregating. We demonstrate that both estimators can be asymptotically biased in the presence of public noise. To address these biases, we suggest instrumenting forecast revisions with past forecast errors. The Anderson-Rubin likelihood ratio test can be applied to test for forecast efficiency, and remains robust even in the presence of weak instrumental variables. Applications of the tests to the U.S. Survey of Professional forecasters covering 1968Q4 through 2016Q4 clearly reveal experts' underreaction to news in their macroeconomic expectations.

Big Data-Driven Economic Index Construction and Tourism Demand Forecasting

Presenter: Wanyi Cao (School of Business, Guangxi University)
Author(s): Wanyi Cao;Mingming Hu;Gang Li;Doris Chenguang Wu

Tourism demand forecasting is vital for destination management. Traditional studies, grounded in neoclassical economic theory, rely on low-frequency macroeconomic variables (e.g., the income level of a country of origin and relative prices at a destination)to model tourism demand. However, this conventional framework is hindered by two limitations: (1) low-frequency data fail to capture real-time market dynamics; (2) a reliance on historical information of variables cannot reflect public expectations of future economic trends. Although big data have played an important role in tourism demand forecasting, critical theoretical challenges persist: How can big data variables be effectively integrated into traditional tourism demand theory? How can economic indices derived from big data be constructed and utilized for tourism demand modeling and forecasting? This study integrates text mining with

domain-specific sentiment lexicon to establish a tourism demand model enhanced by economic indices. Firstly, high-frequency economic indices are extracted from the social media platform (Weibo), leading to the construction of two novel indices: the Income Expectation Index (IEI) and the Relative Price Index (RPI). These indices are then incorporated into SARIMAX and SARIMA-MIDAS models to forecast monthly tourism demand from Mainland China to nine key destinations: Thailand, Hong Kong, Macao, Singapore, Vietnam, Malaysia, Japan, South Korea, and Indonesia. Empirical results indicate that: (1) Models incorporating IEI and RPI outperform those relying on macroeconomic variables alone or their combination, highlighting the superior substitution advantages of big data-driven economic indices over macroeconomic variables in short-term forecasting. (2) SARIMA-MIDAS models that leverage high-frequency economic indices demonstrate greater forecasting accuracy compared to SARIMAX models based on low-frequency economic indices. (3) Among daily, monthly, and mixed frequency setups, the SARIMA-MIDAS models incorporating weekly IEI and RPI variables exhibit the strongest predictive ability. This study makes two key contributions: (1) It pioneers the integration of big data-driven economic indices into tourism demand forecasting, establishing a valuable connection between traditional economic theory and big data variables. (2) It is the first to develop a big data lexicon for economic variables and introduce IEI and RPI as novel tools to enhance tourism demand forecasting accuracy.

Integration of Big Data and Economic Variables in Tourism Demand Modeling

Presenter: Limei Yang (School of Business, Guangxi University)
Author(s): Limei Yang;Mingming Hu;Gang Li

Tourism prices and tourists' income have been identified as primary determinants in modeling tourism demand. Numerous studies have shown that income and price elasticities vary for multiple reasons. Searching for information on the internet has become an indispensable part of tourists' pre-trip planning, marking a significant shift in how travel decisions are made. Serving as crucial bridges between users and vast amounts of information, search engines not only enable users to quickly locate necessary details but also offer personalized recommendations, price comparison tools, and other functionalities by analyzing users' search histories and preferences. The availability of detailed information on the internet alters how income influences tourism demand. Tourists from various income brackets can find travel options that fit their budgets, suggesting that information searches help mitigate the impact of income disparities on tourism demand. Moreover, these online sources also modify how prices affect tourism demand. For instance, by providing tailored recommendations for tourism products, search engines can reshape tourists' perceptions of travel prices, making them more receptive to certain price points. Therefore, the influences of tourism prices and tourists' income on tourism demand vary with the intensity of tourists' online information search behavior. However, how the effects of tourists' income and tourism prices on tourism demand are influenced by online information search behaviors has not been investigated yet. To address the research gaps, this study aims to test the moderating effects of tourists' online information search behavior on tourists' income and tourism prices. An empirical study on tourism demand from eight English-speaking countries to Hong Kong indicates that: 1) Tourists' online information search behavior mitigates the effect of tourists' income on tourism demand; 2) It also mitigates the impact of tourism prices on tourism demand; 3) Online searches related to dining, lodging, shopping, transportation, tours, and recreation consistently confirm these moderating effects. This study enriches the theories of tourism demand in the big data era and provides new insights for governments and tourism administrators regarding online information sharing to enhance tourism marketing.

Progressive Low-rank Multimodal Fusion for Tourism Demand Forecasting

Presenter: Shiqi Liu (Sun Yat-sen University)
Author(s): Shiqi Liu;Chenguang Wu

The integration of heterogeneous data from various sources has demonstrated the ability to enhance the performance of tourism forecasting models, thus receiving increasing attention in recent years. A critical challenge in this domain is information fusion, as existing approaches primarily rely on multivariate time series econometric models that fail to account for inter-modal interactions amongst inputs. Therefore, this study proposes a novel progressive Low-rank Multimodal Fusion (Pro-LMF) model capable of capturing complex nonlinear relationships between multi-source heterogeneous data. The Pro-LMF model not only integrates multimodal inputs with distinct

characteristics into a unified common representation, but also provides fused outputs back to unimodal feature generators through backward connections, enhancing the expressiveness of the fused representations. Daily tourism demand data from mainland China to Hong Kong and Macau spanning January 2023 to November 2024 was used for empirical examination. The results demonstrate that our proposed model significantly outperforms all 4 benchmark models (Naive, ETS, SARIMA and SARIMAX), and further outperforms the static LMF model, confirming its superior performance in tourism demand forecasting.

A Novel Forecast Combination Strategy Based on Marginal Contribution of Predictors

Presenter: Gong Xue (Nanjing University of Science and Technology)
Author(s): Xue Gong;Yudong Wang;Aoran Hong

In the context of the rich-data world, forecasting financial market movements using a wide range of predictors is both common and crucial. This study makes a significant contribution to the literature by proposing a novel forecast combination method that integrates both in-sample and pseudo-out-of-sample information. Through an empirical analysis of behavioral factors in predicting gold futures volatility, our results show that the proposed model outperforms several existing forecast combination methods. These findings are robust across various sensitivity tests. Moreover, our analysis highlights that, compared to investor sentiment and economic policy uncertainty, technical indicators are the most influential behavioral predictors of volatility. Finally, we demonstrate that the proposed model can effectively improve portfolio performance by optimizing risk control.

The Outperformance of High-Frequency Return Forecasted MMVaR in Portfolio Management

Presenter: Li Jiayi (中国科学院大学)
Author(s): Jiayi Li;Zhengjun Zhang

The market application of MMVaR is still in its early stages. While its advantages in portfolio optimization using monthly data have been validated, research on its performance with daily data remains unexplored. This study systematically evaluates the performance of five risk measures in portfolio selection using daily data from the Chinese market and two mature markets. Innovatively incorporating high-frequency data at different time scales to forecast daily MMVaR risk, our findings indicate that when market trends do not reverse, MMVaR demonstrates robust risk measurement capabilities and outperforms traditional risk measures across different market environments. Furthermore, under non-reversing market trends, various risk-return trade-off models exhibit systematic patterns: returns generally follow an inverted U-shape, while standard deviations exhibit a U-shape as confidence levels vary, suggesting that portfolios perform optimally at moderate confidence levels. Integrating both daily and monthly data analyses, MMVaR not only achieves a balance between regulatory requirements and investors' return objectives but also has the potential to become a new benchmark risk measure in the market. This study further expands the academic understanding of MMVaR and provides new insights into its application under real-world trading constraints.

Analyzing Exchange Rate Dynamics within the Global Financial Cycle: A DCC-Copula approach

Presenter: Luis Melo (Banco de la Republica (Central Bank of Colombia))
Author(s): Luis Melo;Jose Romero;Diego Niño

The literature on the Global Financial Cycle, often characterized by fluctuations in international capital flows, asset prices, and risk appetite, has garnered significant attention in recent years. This study employs Dynamic Conditional Correlation Copula models to examine the association between exchange rates and the VIX for a group of seven developed economies and seventeen emerging market economies. This approach provides insights into how exchange rates time-varying correlations behave versus variables related to the Global Financial Cycle, namely the VIX, which influence exchange rates over time and during stress episodes. The findings contribute to understanding further the interconnectedness between time-varying international financial conditions and currency markets, offering valuable implications for policymakers and investors.

Uncertainty Estimation for High-dimensional Nonparametric Forecasts

Presenter: Nuwani Kodikara Palihawadana (Monash University, Australia)
Author(s): Nuwani Kodikara Palihawadana;Rob J Hyndman;Xiaoqian Wang

Point forecasts alone are inadequate unless supplemented by an assessment of their inherent uncertainty. The most common approach to quantifying this uncertainty is through prediction intervals. In this paper, we focus on nonparametric additive index models for high-dimensional time series forecasting, which make no distributional assumptions during model estimation and often produce serially correlated forecast errors, making theoretical prediction intervals infeasible. To address this, we introduce a novel methodology for constructing prediction intervals, called Conformal Bootstrap (CB), which naturally integrates bootstrap with classical split conformal prediction (SCP). Bootstrapping involves resampling in-sample residuals to simulate future sample paths. This use of in-sample residuals hinders the ability to account for the uncertainty introduced by forecasting predictors in ex-ante forecasting scenarios. In contrast, SCP uses out-of-sample non-conformity scores (or forecast errors) obtained from a calibration set, which does not overlap with the training set, to derive prediction intervals for the test set. This use of out-of-sample forecast errors allows the capture of the uncertainty introduced in ex-ante forecasting, assuming future forecast errors for both the forecast variable and the predictors are similar to past errors. However, accounting for the autocorrelation in multi-step forecasts is problematic in SCP. The proposed CB method involves generating out-of-sample forecast errors through sequential splitting of the data, as in SCP. These out-of-sample forecast errors are resampled through a bootstrap procedure, where multi-step forecast errors corresponding to each time point in the calibration set are treated as a single data point. This helps to preserve the temporal dependencies inherent in multi-step forecasts when constructing the prediction intervals. Thus, CB accounts for both the temporal correlation in multi-step forecasts and any additional uncertainty introduced by ex-ante forecasting, leveraging the strengths of both bootstrap and SCP. The CB prediction intervals are evaluated against block bootstrap prediction intervals and those derived from existing conformal prediction methodologies within the generic online learning framework proposed by Wang & Hyndman (2024), through empirical applications to forecast heat-exposure-related daily mortality and daily solar intensity.

Probabilistic forecasting of intermittent time series: local vs global models

Presenter: Stefano Damato (IDSIA USI-SUPSI)
Author(s): Stefano Damato;Nicoló Rubattu;Dario Azzimonti;Giorgio Corani

On smooth time series, even simple global methods can outperform state-of-the-art local models. Instead, most models for forecasting intermittent time series are still local. We compare local and global probabilistic models for intermittent time series. Our study includes state-of-the-art local (iETS, TweedieGP) and global models (DeepAR and Autoformer). We also consider simpler models like feed-forward neural networks. We discuss the choice of distribution head for neural models by comparing the negative binomial and the Tweedie distribution. Indeed, we apply for the first time a fully evaluated Tweedie distribution to deep neural networks. We present results based on thousands of time series. Global models often present unstable results, particularly when datasets have few series or these are short. This shortcoming, together with the higher computational cost, makes the use of local models often preferable.

A Prototype Model for State Revenue Forecasting under High Fiscal Uncertainty

Presenter: Cheng Yang (Liaoning University)
Author(s): Cheng Yang;Kajal Lahiri

Tax revenues have become increasingly volatile and more challenging to forecast as they are affected not only by business cycles, but also by newly enacted policy changes and uncertain federal grants. We forecast yearly tax revenues and their uncertainties for New York at monthly frequencies at horizons from one to 18 months by combining three predictive blocks: a model with monthly tax revenues, a model with US factor and a model with New York factor with large number of predictors, where the factors are estimated jointly by multi-region dynamic factor models. We augment each of these models by a novel tax policy variable to incorporate expected revenue changes due to new policies using executive budget reports and fiscal surveys. By augmenting the tax policy variable to mixed-frequency models, we can improve the forecast accuracy by up to 32.8% over models without it. The tax

policy variable captures expectations of government budget experts that are not available in published data. We use a new blocking-based-residual-bootstrap method to estimate uncertainties that are updated monthly together with our combined forecasts. The uncertainties are found to increase with the length of the forecast horizon in general and become large due to recessions and shocks like COVID-19. Finally, the forecasts and their uncertainties serve as a forward-looking way to help set rainy day funds (RDF) in the context of increasing tax revenue volatility.

Carbon Emission Prediction in Rural Areas of China: Visual Machine Learning Models

Presenter: Hejing Wang (Beijing Institute of Technology)
Author(s): Hejing Wang;Hejing Wang

The carbon emissions in rural areas of China have steadily increased and now account for about 15% of the country's total carbon emissions. Therefore, accurately predicting rural carbon emissions has become a critical task for the government to formulate effective plans and achieve the dual carbon goals. However, existing research has overlooked the nonlinear characteristics of rural carbon emissions and has failed to incorporate significant influencing factors such as rural labor transfer. In view of this, this study employs machine learning methods capable of learning and mining intrinsic patterns from data, aiming to predict rural carbon emissions more accurately and effectively. Specifically, we first identified 15 key explanatory features based on historical rural carbon emission data from 30 Chinese provinces during 2006-2022, utilizing feature selection algorithms. Subsequently, the XGboost model was identified as demonstrating superior predictive performance among seven machine learning models evaluated using five performance indicators.

Air Quality Index Prediction Based on Spatio-Temporal Graph Neural Networks

Presenter: Yifan Yang (Xi'an Jiaotong University)
Author(s): Yifan Yang;Ju'e Guo

Air pollution, particularly as measured by the Air Quality Index (AQI), poses significant threats to public health and the environment, necessitating accurate prediction models for timely interventions. This study introduces a novel Spatio-Temporal Graph Neural Network (STGNN) model designed to predict AQI across multiple monitoring stations. The model leverages spatial and temporal data, employing various time-series distance metrics to construct an adjacency matrix that captures spatial dependencies between stations. By integrating these complex spatio-temporal relationships, the STGNN model significantly outperforms traditional methods in prediction accuracy. This study not only provides a powerful tool for air quality prediction in Xi'an but also offers a scalable framework applicable to other regions and varying environmental conditions. The importance of accurate AQI prediction is underscored by the severe health impacts of air pollution, particularly PM2.5. Traditional prediction models, such as Numerical Weather Prediction (NWP) and ARIMA, face limitations in computational efficiency and their ability to capture complex nonlinear characteristics and spatial information. Graph Neural Networks (GNNs) have emerged as a promising approach due to their ability to handle complex network data and capture spatial relationships. Spatio-Temporal Graph Neural Networks (STGNNs) extend this capability by incorporating temporal dynamics, making them highly effective for tasks like traffic flow prediction and electricity price forecasting. Despite their potential, the application of STGNNs to AQI prediction remains underexplored. This study addresses this gap by developing an STGNN-based model that not only preserves the temporal trajectory of AQI data but also exploits interdependencies among monitoring stations, thereby enhancing prediction accuracy. Key innovations of this study include the integration of spatio-temporal information through STGNN, a novel method for constructing adjacency matrices using multiple time-series distance metrics, and comprehensive experimental validation. The proposed framework demonstrates exceptional predictive ability across different regions and time periods, highlighting its broad application potential for air quality management. This research contributes a robust and scalable tool for accurate AQI prediction, paving the way for more effective environmental management strategies.

Interpretable Daily Carbon Emission Forecasting in China's Transport Sector: Application of Temporal Fusion Transformer

Presenter: Ming Yang (Dalian Maritime University)
Author(s): Ming Yang;Haiying Liu

Backgrounds: The transportation sector in China, contributing over 12% of Chinese overall CO_2 emissions, presents urgent decarbonization challenges. Accurate daily carbon emission forecasting is crucial for identifying carbon peaking timelines and formulating evidence-based mitigation strategies. However, existing prediction models predominantly focus on univariate time-series analysis, overlooking critical external determinants. Methodologies: This study proposes an interpretable deep learning framework using Temporal Fusion Transformer (TFT) to enhance prediction accuracy while revealing key influencing mechanisms. We analyze daily emission data (June 19, 2021-May 17, 2024) from Carbon Monitor covering domestic aviation, international aviation, and ground transportation. Feature variables include: Energy market indicators: WTI crude closing prices and carbon trading turnover prices. Public attention: related Baidu search indices. Temporal features: time stamps and calendar effects with holiday lags/leads. COVID-19 mobility restrictions. The model architecture employs direct multi-step forecasting with a 14-day lookback window to predict emissions over the subsequent 4-day horizon. Results: Experimental results demonstrate TFT's superior performance over baseline models in both in-sample and out-of-sample tests. Interpretability analysis reveals the following dominant factors in descending order: Statistic variables: different sectors. Encoder variables: WTI crude closing prices, historical carbon emission, and lag-holiday effects (7-day lag). Decoder variables: influence of COVID-19, time index, and day of week. Conclusions: The application of TFT addresses mitigation of error propagation by replacing iterative multi-step forecasting with direct sequence-to-sequence prediction and recognizing and processing features adaptively. Practical implementation of TFT enables policymakers to formulate data-driven decarbonization strategies and enhance capacity to energy price volatility, public supervision, and uncertainty events.

Feature Selection for Neural Network Forecasting - an Empirical Evaluation of Partial Dependence, Pertubation and Gradient Based Techniques for Railway Revenue Forecasting

Presenter: Sven F. Crone (iqast - ai forecasting software // Lancaster University)
Author(s): Sven F. Crone

Feature selection is considered of preeminent importance for specifying accurate, robust and efficient machine learning methods, attracting over 22,000 papers for specifying popular machine learning (ML) methods. While some prominent algorithms include feature selection explicitly in their methodology, such as tree-based methods of random forests or xgboost by limiting the depth of base learners while prioritising splits on features with highest feature importance. For time series data, most papers employ only simple heuristics such as p= n*s autoregressive lags yt-p, with s = seasonal length, and n typically set between 1 and 3 (see also the nnetar function of the popular forecast package in R, Hyndman and Caceres, 2024). Crone and Kourentzes utilise established statistical techniques, such as ACF and PACF, as well as Stepwise regression, for neural network feature selection. Zimmermann suggest to unfold partial dependence plots across time series and check importance, consistency and nonlinearity (Zimmermann, 2020). More recent methods determine feature importance based on the predictive performance of the output, which include partial dependence plots, pertubation baeed feature importance, both of which are applicable for feature selection across ML methods, and gradient based feature importance utilising the learning information of neural networks during parameterisation. However, despite their prominence in classification, clustering and regression research, only few papers consider these feature selection methods on time series data for forecasting. This paper seeks to address this gap by comparing statistical approaches and heuristics with three feature selection approaches from ML: pertubation, partial dependence and gradient weight based approaches to determine the empirical accuracy on a real world railway demand dataset with many features. The results indicate the promise of more advanced approaches of feature importance over simpler methods, and find computationally efficient trade-offs between accuracy and speed of feature selection.

Active Learning Markets for Forecasting: How to Purchase the Most Informative Data

Presenter: Xiwen Huang (Technical University of Munich)
Author(s): Xiwen Huang;Pierre Pinson

This work presents an innovative active learning market framework tailored for forecasting applications in resource-constrained environments. Unlike traditional forecasting models that assume bulk or unlimited data availability, our approach selectively acquires only the most informative data, balancing prediction accuracy and cost efficiency. We implement active learning in batch settings to optimize one-time data purchases and extend the framework to an online setting where real-time adaptive data acquisition is required. Our methodology integrates market-based optimization for data selection and exponential forgetting for efficient model adaptation. Empirical validation on energy forecasting and marketing prediction tasks shows that our approach outperforms benchmark random sampling, leading to lower data costs and improved forecasting accuracy. Furthermore, our framework is scalable to multi-buyer and multi-seller settings, making it highly adaptable across energy, finance, and industrial forecasting applications. This research provides a cost-effective, adaptable strategy for optimizing data acquisition in forecasting models, contributing to the broader field of intelligent forecasting and decision-making.

PatchLinear: Patching Based Linear Model for General Time Series Analysis

Presenter: Santosh Palaskar (Unilever)
Author(s): Santosh Palaskar;Nandyala Hemachandra;Narayan Rangaraj

Time series analysis has extensive applications across diverse domains, including supply chain management, electricity consumption, healthcare, and sports. These applications often rely on tasks such as forecasting, anomaly detection, action recognition, and classification. Recent advancements in the field have predominantly focused on transformer-based architectures, which require large datasets and significant computational resources. These models are very effective for sequence-based predictions due to their permutation invariance property. However, they often lose critical temporal information in time series data due to the limitations of channel mixing and attention mechanisms. In this work, we argue that a linear model is sufficient to capture and extract essential temporal information from time series data. To support this claim, We propose a novel patch-based linear model, PatchLinear, that captures local semantics using patch-based embeddings and leverages weight sharing across channels to model inter-channel influence, enabling the learning of high-level representations. PatchLinear demonstrates superior performance over existing transformer-based and traditional statistical models across various time series analysis tasks like forecasting, anomaly detection and classification. Our lightweight model, with 127.7K parameters (compared to DLinear's 139.7K parameters and transformer-based model with over a million parameters), offers a computationally efficient and robust solution for time series analysis.

ARTree: A Deep Autoregressive Model for Phylogenetic Inference

Presenter: Tianyu Xie (Peking University)
Author(s): Tianyu Xie;Cheng Zhang

Designing flexible probabilistic models over tree topologies is important for developing efficient phylogenetic inference methods. To do that, previous works often leverage the similarity of tree topologies via hand-engineered heuristic features which would require domain expertise and may suffer from limited approximation capability. In this paper, we propose a deep autoregressive model for phylogenetic inference based on graph neural networks (GNNs), called ARTree. By decomposing a tree topology into a sequence of leaf node addition operations and modeling the involved conditional distributions based on learnable topological features via GNNs, ARTree can provide a rich family of distributions over tree topologies that have simple sampling algorithms, without using heuristic features. We demonstrate the effectiveness and efficiency of our method on a benchmark of challenging real data tree topology density estimation and variational Bayesian phylogenetic inference problems.

High-dimensional Quantile Vector Autoregression with Influencers and Communities

Presenter: Huang Shiwei (School of Management, University of Science and Technology of China)
Author(s): Shiwei Huang;Yu Chen

Quantile vector autoregression (QVAR) models offer enhanced capabilities over vector autoregression (VAR) models in analyzing asymmetric interactions within multiple time series and are widely used in many areas, including finance and economics. However, in scenarios involving high-dimensional data where the number of time series N exceeds the length of the time series T, both parameter estimation and the theoretical establishment present significant challenges. Notably, existing research has not yet explored these challenges with the framework of QVAR models. To handle this problem, we propose a novel high-dimensional QVAR model that incorporates influencers and communities, which assumes that variables within distinct communities have shared dependency structures and are influenced by the same set of variables, called influencers. We develop an estimation procedure based on the alternating minimization algorithm and the convolution-smoothed approach. The local consistency results for the estimated parameters are established in high-dimensional settings with sub-Weibull innovations. Numerical studies illustrate that the proposed model performs well in finite samples. The proposed model is applied to identify the most influential macroeconomic variables in the United States across business cycles and construct a dynamic quantile network with influencers and communities for volatility spillover effects of global stock markets.

Merging VAR and factor model with One-sided Dynamic Autoregressive Principal Components

Presenter: Yangzhuoran Fin Yang (Maastricht University)
Author(s): Yangzhuoran Fin Yang

This paper proposes a multivariate model structure to forecast high-dimensional time series. The model extracts components as linear combinations of the time series present and past values. The components follow autoregressive processes and are called One-sided Dynamic Autoregressive Principal Components. They contain signals useful for forecasting that are shared across time series. The original time series are then reconstructed from the fitted component values. This model contains three stages of transmission of information: from the time series to the components, from the components to the components' future, and from the components back to the original time series. The effects of these stages are captured by different sets of parameters, which are estimated by minimising reconstruction error. Special cases of this general specification cover a range of models including constrained Vector Autoregression and factor models. Forecasting performance is examined using simulation and empirical applications.

Bayesian VAR Models with a Combination of the DSGE Model-Implied Prior and the SSVS in Mean-IW Prior

Presenter: Xin Zheng (Beijing Institute of Technology)
Author(s): Xin Zheng

This study combines the Dynamic Stochastic General Equilibrium (DSGE) prior and the Stochastic Search Variable Selection (SSVS) in mean-Inverse Wishart (IW) prior into a hybrid prior to estimate Bayesian Vector Autoregression (BVAR) models. This hybrid prior innovatively shrinks unimportant VAR coefficients toward zero but centers important VAR coefficients on the DSGE model-implied prior means, while preserving error covariance stability through the IW prior. Based on root mean square forecast errors, predictive log likelihoods, and Bayesian model comparison, our main findings are two-fold. For short-term forecasting, BVAR models with the Minnesota prior significantly outperform Bayesian DSGE-VAR models with the Normal-Inverse Wishart (N-IW) prior and the SSVS in mean-IW prior, at a 5% significance level. For long-term forecasting, Bayesian DSGE-VAR models with the SSVS in mean-IW prior significantly outperform BVAR models with the Minnesota prior and the N-IW prior, at a 5% significance level. The combination of the DSGE prior and the SSVS in mean-IW prior in BVAR models embodies theoretical structural coherence from the DSGE model-implied restrictions and empirical shrinkage efficiency from the SSVS in mean-IW technique. These findings contribute to policy makers in forecasting macroeconomic series over a long horizon.

Comparison in Dynamic Forecasting: A Bayesian LASSO State-Space and Bayesian Factor VAR Analysis

Presenter: Mingchuan Zhou (University College Dublin)
Author(s): Mingchuan Zhou

This study develops a Bayesian LASSO state-space model to forecast key macroeconomic variables—including the Federal Funds Rate (FEDFUNDS), Industrial Production Index (INDPRO), and Consumer Price Index (CPI)—using a macro-finance dataset by McCracken and Ng from the Research Division at the Federal Reserve Bank of St. Louis. The model employs Kalman filtering for latent state updates and Bayesian LASSO regularization via Gibbs sampling, forming a dynamic mechanism that stabilizes forecasts and improves accuracy under uncertainty from various financial factors. Empirical analysis demonstrates that the Bayesian LASSO state-space model outperforms a benchmark moving average model in out-of-sample forecasts within financial datasets, specially, forecasts of FEDFUNDS and INDPRO exhibit significantly improved accuracy in both point and density forecasting, Mean Squared Error of 0.06 and Root Mean Squared Error of 0.24, and log scores of -0.41 and -0.96, respectively. These results indicate that the model effectively captures the dynamic patterns in financial time series, demonstrating sensitivity to the dynamic selection of informative priors and the degree of shrinkage. Additionally, robustness checks on macroeconomic and combined datasets confirm the adaptability of this model to various macroeconomic dynamics. For comparative purposes in machine learning and traditional econometric forecasting via Bayesian methods, a Bayesian Factor-Augmented Vector Autoregression (FAVAR) model with pandemic priors is also introduced to examine exogenous shocks. The FAVAR employs rolling-window out-of-sample forecasts, smoothly updating parameters over time and achieving stable performance across both financial and macroeconomic data, outperforming static models (VAR, BVAR, FAVAR) in pandemic periods. Robustness tests indicate that the Bayesian LASSO state-space model maintains strong predictive performance across diverse macro-financial datasets under a fixed forecasting approach, while the Bayesian FAVAR better captures evolving economic structures through rolling-window predictions. By integrating data-driven Bayesian machine learning with theory-driven structural econometrics, our proposed approach balances statistical rigor with economic intuition, offering robust theoretical and empirical support for dynamic forecasting and parameter estimation in macroeconomic research.

Big Model-Based GDP Forecasting: The Role of Micro-Level Accounting Information

Presenter: Harold Guo (Center for Forecasting Science, Chinese Academy of Sciences)
Author(s): Harold Guo;Yongmiao Hong

Accurate economic forecasting constitutes a critical foundation for evidence-based policy formulation and implementation. This study adopts a machine learning framework through the lens of large-scale models (big models) to incorporate heterogeneous micro-level accounting information across economic entities for macroeconomic prediction. Five principal findings emerge: Firstly, Integration of firm-specific indicators (e.g., Return on Assets, ROA) significantly enhances year-on-year GDP forecast accuracy, demonstrating 25.7% and 57.7% reductions in RMSE for short-term and long-term predictions respectively. These improvements persist across robustness checks involving temporal windows, sample variations, and initialization dates. Feature importance analysis reveals the model's adaptive feature selection mechanism as the enhancement driver. Secondly, among 20 examined accounting factors, only profitability-related metrics exhibit statistically significant predictive power. Thirdly, heterogeneity simulations using synthetic data demonstrate that microdata superiority originates from firm-specific variations. The RMSE hierarchy across information types follows: microdata < joint distribution < independent distribution < aggregated data, with maximum RMSE reduction of 63.3% compared to aggregated data. Fourthly, the predictive advantage of microdata generalizes to quarter-on-quarter GDP forecasting. Fifthly, cross-validation between macroeconomic price variables and micro-level profitability factors achieves complementary information integration. The combined model attains a minimum RMSE of 0.87, representing 36.1% improvement over benchmark models.

Expected Return in Night and Day: Role of Trading Volume

Presenter: Zhiheng He (Tsinghua University)
Author(s): Zhiheng He;Alfred Qi Fan;Dashan Huang

We document a novel high-volume overnight premium (intraday discount). Specifically, stocks with high trading volume exhibit remarkable future outperformance (underperformance) during the overnight (intraday) period. This phenomenon is also prevalent in global equity markets, robust to alternative measures of trading volume, and persistent for at least three years after portfolio formation. In a tractable equilibrium model, we argue such heterogeneous night/day effect is the result of the cycling pattern of belief dispersion being high/low at market open/close, due to the overconfidence of late-informed investors. To justify this idea, we use high-frequency returns and document the decreasing cross-sectional dispersion of market betas over the intraday period is especially strong for high-volume stocks. We also propose several testable model predictions and empirically shows that our main results are more pronounced when uninformed investors are more optimistic, or more dominant in overall trading activity.

Sparsity-Induced Global Matrix Autoregressive Model with Auxiliary Network Data

Presenter: Dan Yang (The University of Hong Kong)
Author(s): Dan Yang;Sanyou Wu;Yan Xu;Long Feng

Jointly modeling and forecasting economic and financial variables across a large set of countries has long been a significant challenge. Two primary approaches have been utilized to address this issue: the vector autoregressive model with exogenous variables (VARX) and the matrix autoregression (MAR). The VARX model captures domestic dependencies, but treats variables exogenous to represent global factors driven by international trade. In contrast, the MAR model simultaneously considers variables from multiple countries but ignores the trade network. In this paper, we propose an extension of the MAR model that achieves these two aims at once, i.e., studying both international dependencies and the impact of the trade network on the global economy. Additionally, we introduce a sparse component to the model to differentiate between systematic and idiosyncratic cross-predictability. To estimate the model parameters, we propose both a likelihood estimation method and a bias-corrected alternating minimization version. We provide theoretical and empirical analyses of the model's properties, alongside presenting intriguing economic insights derived from our findings.

On Semiparametric Real-time GARCH Models

Presenter: Muyi Li (Xiamen University)
Author(s): Muyi Li

We propose a semiparametric real-time GARCH (SRT-GARCH) model for volatility forecasting under nonstationary market conditions. The model decomposes return volatility into a smooth time-varying unconditional component and a short-run GARCH-type innovation structure augmented with current-period information. The long-run component is estimated nonparametrically using kernel methods, while the remaining parameters are estimated via conditional maximum likelihood. We establish the consistency and asymptotic normality of the estimators, with the parametric rate achieved in the second step. Monte Carlo simulations demonstrate the favorable finite-sample performance of the proposed SRT-GARCH models. Empirical applications illustrate that the proposed SRT-GARCH models yields significantly improved forecasting accuracy and Value-at-Risk estimates,compared to the existing nonstationaty GARCH models.

A Study for Forecasting Solar Energy Power Generation.

Presenter: Young Doo Jeong (Department of Statistics and Data Science, University of Seoul)
Author(s): Young Doo Jeong;Arnie de Castro;Tae Yoon Lee;Hae Rim Kim;Dae Hyuk You;Hyeonmin Jeon

Due to Korea's high dependence on energy imports, it is important to balancedomestic energy production and consumption. As a result, accurate prediction of the amount of solar power is becoming more important as the supply of renewable energy increases, centering on solar power in Korea. Since solar energy output fluctuates according to weather and climate changes, an accurate prediction modelis needed to increase the stability of the electrical grid and prevent an imbalance in supply and demand. In this study, a convolutional neuralnetwork (CNN), a short- and

long-term memory network (LSTM), and a transformer model were combined to study an effective deep learning framework for photovoltaic energy prediction. The CNN-LSTM-Transformer hybrid model effectively models short-term fluctuations and long-term trends by capturing both spatial and temporal dependencies of photovoltaic data. In addition, a time-diffusion model has emerged that can increase the performance of time series prediction models by using synthetic data that can be used when there is little training data. by combining the synthetic data, the generalization ability and prediction accuracy of the model were improved. It is shown that the CNN-LSTM-Transformer perform well in accurately predicting the solar energy output in various prediction periods.

Nationwide Weather Impact on Photovoltaic Generation Forecast

Presenter: Jae in Song (Korea Power Exchange)
Author(s): Jae In Song;Hyunsu Kim;Seok Min Choi;Yo-Hwan Choi;Kun Suk Lee;Chang Gun Lee

Accurate forecast of nationwide photovoltaic (PV) generation is crucial to electric power system operation. It is well known that weather variables are highly related to the PV generation of individual PV panels. However, directly utilizing this relationship for nationwide PV generation forecasting remains challenging due to spatially and temporally limited weather observations and the diverse characteristics of PV installations across regions. To overcome such limits, we investigated the relationships between nationwide weather variables and PV generation. At the conference, three key findings of this study will be presented. First, we discovered that the nationwide weighted average of weather variables exhibits a strong correlation with PV generation. Next, our analysis indicates that air temperature significantly influences PV generation, ranking second only to solar irradiance. Lastly, satellite imagery can be effectively utilized to estimation of solar irradiance on a nationwide scale. Detailed methodology used in this study also will be introduced at the presentation. Keywords: PV generation, Weather, KPX, Multivariate linear regression, satellite imagery

Development of Operational PV Generation Forecast Model in Korea Based on CNN and K-means Clustering Method

Presenter: Chae Rin Kim (KPX(Korea Power Exchange))
Author(s): Chae Rin Kim;Kun Suk Lee;Heung-Gu Son;Minho Song;Chang Gun Lee

Photovoltaic (PV) power generation significantly influences overall electricity demand by reducing net load; therefore, accurately forecasting PV generation is essential for precise power demand prediction. To address this, the Korea Power Exchange (KPX) has developed and implemented a specialized PV forecasting model called the KPX Irradiance and Photovoltaic Forecasting System (KIFS). The KIFS model integrates Convolutional Neural Networks (CNN) with K-means clustering techniques to forecast PV generation for 40 regions across Korea, providing predictions up to 10 days ahead. The forecasting process is conducted in two sequential steps: first, regional solar irradiance is predicted, and subsequently, regional PV generation outputs are estimated based on these irradiance forecasts. For the solar irradiance prediction component, KIFS utilizes six key meteorological variables provided by the Korea Meteorological Administration (KMA), those are temperature, humidity, cloud cover, wind speed, precipitation, and top-of-atmosphere solar radiation. Given the variability of meteorological conditions across different regions, KIFS performs localized predictions of solar irradiance individually tailored for each of the 40 regions. Finally, these irradiance predictions serve as critical input data for accurately forecasting regional PV generation. Detailed methodologies, technical implementations, and performance evaluations of the KIFS model will be thoroughly presented during the forthcoming conference. Keywords: PV generation forecast, power demand, KPX, KMA, k-means clustering, CNN(Convolutional Neural Network)

News-Driven Load Forecasting: Generative Agents and Large Language Models for Unstructured Data and Event Analysis

Presenter: Xinlei Wang (Univerisity of Sydney)
Author(s): Xinlei Wang;Jinjin Gu;Jing Qiu;Guolong Liu;Junhua Zhao;Zhaoyang Dong

This paper proposes a novel approach, Intelligent Text-Analytic Load Forecasting (ITALF), to address the short-term

load forecasting (STLF) problem by utilizing Large Language Models (LLMs) and Generative Agents. It emphasizes the challenges faced by traditional forecasting methods in adapting to rapid changes and complex patterns in energy consumption, particularly during unexpected social events. Our method leverages LLMs to process and analyze a mix of unstructured data sources, including historical load data, news articles, calendar dates, and weather forecasts, thereby enhancing the accuracy and adaptability of STLF. An LLM-based agent with reasoning capabilities is introduced to select and understand relevant news, demonstrating the model's ability to integrate diverse information for more precise forecasting. A reflection agent is also introduced to evaluate the prediction results and refine the news selection logic. Our prediction results surpass all baseline models in predictive accuracy, indicating that LLMs excel in managing the complexities of load forecasting patterns. This innovative approach not only improves forecasting accuracy but also indicates potential shifts in STLF paradigms by integrating unstructured data through advanced AI techniques.

Dual Multi-objective Heuristic Optimization Driven Kernel Ensemble for Oil Price Forecasting

Presenter: Zhaorong Huang (Sichuan University)
Author(s): Zhaorong Huang;Lean Yu;Bo Yang

In the field of crude oil time series forecasting, despite the surging popularity of heuristic optimization algorithms, machine learning methods, and ensemble models, their integrated applications have yet to grapple with critical challenges. The discussion on how to combine interval uncertainty with other modules to enhance point prediction is insufficient, since characterized by its reliance on the meticulous adjustment of a uniform feature system and model structure, defines the construction of upper-lower bounds or multiple confidence intervals. However, this static paradigm falls into the single parameter mode trap and the simplification fallacy of uncertainty propagation paths. Consequently, it struggles to capture the dynamic nature of time-varying feature and the intricate interaction effects among subsystems, which leads to an embarrassed situation where point prediction results across diverse scenarios do not benefit substantially from the description of interval uncertainty. To tackle these obstacles, this paper presents a dual multi-objective heuristic optimization-driven kernel ensemble framework that systematically integrates multiple heuristic applications. The proposed approach simultaneously optimizes four key aspects: feature selection, sub-model architecture design, ensemble sub-model selection, and interval learning, to reconcile the conflicting requirements of structural diversity and data specificity in kernel ensemble modeling. Specifically, a prediction interval uncertainty-based generation strategy is employed to induce differentiated hybrid kernel network structures, forming a Pareto set of representative point prediction kernel models. This two-stage kernel ensemble algorithm from both kernel function ensemble and kernel regression ensemble perspectives constrains the evolutionary trajectory of kernel models represented by heuristic population individuals, emphasizing the cohesive integration of these kernel structure that maximizes diversity. Empirical experiments involving total of 43 comparisons with common feature selection methods, machine learning algorithms, and ensemble strategies, as well as ablation studies, demonstrate that the proposed framework achieves superior performance and enhanced robustness in high-dimensional multi-variable monthly and weekly oil price forecasting.

T2f: Integrating Actor-Critic Reinforcement Learning with Temporal Attention for Local-Global Time-series Forecasting

Presenter: João Filipe Pinto Sousa (NOVA Information Management School (NOVA IMS))
Author(s): João Filipe Pinto Sousa;Roberto Henriques

Time-series forecasting of multiple related sequences presents unique challenges due to the complex interplay between individual series characteristics and global patterns. We present T2f, a novel forecasting method combining ensemble learning with an actor-critic architecture based on the Twin Delayed Deep Deterministic (TD3) algorithm. T2f addresses the local-global pattern recognition tension through a two-stage architecture. The local stage employs multiple complementary models (VAR, LSTM, GRU, and adaptive sequential models) to capture series-specific dynamics. The global stage uses TD3-based reinforcement learning to determine optimal model combination weights dynamically, framing forecasting as a sequential decision-making problem. Key innovations include a temporal attention mechanism that adaptively weights historical observations based on recency and volatility, and the

integration of a FEDFormer-inspired layer that exploits sparse frequency domain representations. Unlike traditional methods optimizing solely for statistical accuracy, T2f incorporates context-aware error measurement through a reward function combining symmetric mean absolute percentage error (sMAPE) with weighted symmetric mean absolute error (wsMAE). We evaluate T2f on five diverse datasets: Household Power Consumption, Electricity Load Diagrams, two Electricity Transformer Temperature sets, and a synthetic dataset. Experimental results indicate that T2f reduced mean absolute error by up to 35% compared to statistical models and showed improvements of up to 40% on context-weighted metrics in specific scenarios. While specialized models occasionally outperformed T2f on highly regular patterns, it demonstrated superior adaptability to contextual weights and complex temporal dynamics. The method's consistent performance across diverse forecasting scenarios, coupled with its ability to incorporate domain-specific priorities, makes it particularly valuable for applications where forecast errors have varying operational impacts.

A Zero-Shot LLM for Expert Forecast Combination

Presenter: Yinuo Ren (School of Advanced Interdisciplinary Sciences, University of Chinese Academy of Sciences, 19 Yuquan Road,Beijing, 100049, Beijing, China)
Author(s): Yinuo Ren;Han Feng;Pengyang Song;Hanfeng Zhou

This study explores the potential of large language models (LLMs) to enhance expert forecasting through ensemble learning. Leveraging the European Central Bank's Survey of Professional Forecasters (SPF) dataset, we propose a comprehensive framework to evaluate LLM-driven ensemble predictions under varying conditions, including the intensity of expert disagreement, dynamics of herd behavior, and limitations in attention allocation. The framework incorporates key contextual factors, such as high-volatility periods and phases of elevated expert divergence, to ensure robust and nuanced analysis. Our empirical investigation compares forecast errors between LLM-based methods and traditional equal-weight (EW) techniques. The findings consistently demonstrate that LLMs outperform EW methods, particularly in high-uncertainty environments where expert disagreement is significant or attention signals are sparse. Specifically, the improvement in predictive accuracy is substantial, with a 10% reduction in mean squared error (MSE) compared to conventional approaches. These results underscore the potential of LLMs as advanced tools for improving expert-driven predictions in volatile and uncertain environments. Furthermore, LLMs exhibit remarkable adaptability in few-shot and zero-shot learning paradigms, underscoring their utility in data-scarce forecasting scenarios where conventional methods often struggle. By advancing methodological innovations in forecast combination, this study addresses critical questions about the role of AI in augmenting expert judgment under uncertainty. The results highlight the transformative potential of LLMs to bridge gaps in expert-driven predictions, offering both practical applications for policymakers and theoretical contributions to the field of predictive analytics.

Is VMD Based Decomposition-reconstruction-ensemble Mode Effective for Price Forecasting? Evidence from a Multivariate Multi-scale Ensemble Model

Presenter: Na Li (Sichuan University)
Author(s): Li Na; Zhaorong Huang; Wang Xi; Bo Yang; Lean Yu

In the context of crude oil price forecasting, the application of decomposition-reconstruction ensemble prediction models spans a history of a decade or so, for this prediction framework can clearly analyze the internal structure and changing patterns of data by decomposing oil price series into long-term trends, seasonal fluctuations, and random noise. However, latest research criticizes their effectiveness only valid in long-term horizon after data leakage issues are taken into account, while others argue that decomposition modes possess the highest feature importance at all prediction horizon. This paper extends this complex univariate forecasting debate to the context of multivariate forecasting field by constructing a feature variable set of a hundred of dimensions for WTI and Brent weekly futures prices. At the methodological level, we thoroughly examine and control the negative endpoint effects in real-time decomposition by revising the frequency of the variational mode decomposition algorithm. Our stable decomposition-reconstruction-ensemble prediction system with feature engineering and hierarchical ensemble modules fully explores the synergistic effects of multiple oil influencing factors across different time scales. Comparative experiments with machine learning and variable shrinkage methods show that dimensionality reduction regression based on principal component analysis (PCA) plays a pivotal predictor role on the smooth variational modal components in multivariate

scenarios. The multiscale hierarchical PCA prediction system has smaller horizontal prediction errors and higher directional accuracy and computing efficiency, with significant forecasting advantages compared to recent multi-scale decomposition models. Owing to the linearization of the module in our multivariate multi-scale ensemble model, sound interpretability has been derived by local interpretable model-agnostic explanations (LIME), as a practical tool to unravel the black box of decomposition-based models. Further discussion with different decomposition algorithms affected by endpoint effects and ablation experiments indicate that the decomposition and reconstruction steps only serve as an integrative medium for fusing multi-scale influencing factor information without additional information beyond the timeseries itself. Thus, this paper positively responds the key advantages and limitations of decomposition-based models in a weakly efficient market.

Penalising Forecast Accuracy? Data Valuation Challenges in Power Distribution Systems

Presenter: Andrey Churkin (Imperial College London)
Author(s): Andrey Churkin;Pierre Pinson

Power systems are inherently dependent on the quality of data used to solve forecasting, operational, and planning tasks. Recently, the integration of renewable energy sources, battery energy storage systems, and active demand has led to the emergence of independent data providers who generate valuable forecasts. However, such providers often lack incentives to share private operational energy data and forecasts with third parties. In this regard, data marketplaces offer a promising tool for incentivising data sharing, improving forecasting, and optimising decision-making in power systems. There is ongoing research on developing data valuation mechanisms and market designs for energy data trading. Existing studies propose evaluating the usefulness of aggregate forecasts and compensating data providers based on the cost savings enabled by improved forecasting. This work explores data valuation challenges in distribution networks with independent data providers. It is assumed that electric vehicle aggregators (EVAs) can predict and to some extent control their charging power, thus providing valuable day-ahead load forecasts for the system operator. To estimate the value of EVAs' load forecasts, this work leverages optimisation models for reducing the system's day-ahead dispatch and real-time balancing costs. Data providers are then compensated based on the avoided costs due to their load forecasts. However, such seemingly simple and fair mechanisms may not be directly applicable in realistic power distribution, which must operate under nonlinear constraints. The simulations demonstrate that popular data valuation frameworks, such as Shapley value-based methods, can yield counterintuitive results. For instance, forecast error cancellation can make less accurate forecasts seem more beneficial than precise ones. As a result, data providers may be penalised for sharing accurate forecasts that inadvertently increase system costs. The value of load forecasts may also vary significantly due to locational factors (e.g., connections of charging stations) and network constraints (e.g., voltage limits). These findings expose the complexity of data valuation in non-linear systems, emphasising the need for new mechanisms that ensure practical applicability, incentive compatibility, individual rationality, and fairness in data markets.

Multi-Task Decision-Focused Prediction in Power Systems: Bridging from Renewable Prediction to Voltage Regulation

Presenter: Sang Linwei (Southeast University)
Author(s): Linwei Sang

Prediction plays a vital role in the active distribution network voltage regulation under the high penetration of photovoltaics. Current prediction models aim at minimizing individual prediction errors but overlook their collective impacts on downstream decision-making. Hence, this paper proposes a safety-aware semi-end-to-end coordinated decision model to bridge the gap from the downstream voltage regulation to the upstream multiple prediction models in a coordinated differential way. The semi-end-to-end model maps the input features to the optimal var decisions via prediction, decision-making, and decision-evaluating layers. It leverages the neural network and the second-order cone program (SOCP) to formulate the stochastic PV/load predictions and the var decision-making/evaluating separately. Then the var decision quality is evaluated via the weighted sum of the power loss for economy and the voltage violation penalty for safety, denoted by regulation loss. Based on the regulation loss and prediction errors, this paper proposes the hybrid loss and hybrid stochastic gradient descent algorithm to back-propagate the gradients of the hybrid loss with respect to multiple predictions for enhancing decision quality. Case studies verify the effectiveness

of the proposed model with lower power loss for economy and lower voltage violation rate for safety awareness.

Multi-objective Probabilistic Forecast Combination: Decision Perspective

Presenter: Shengjie Wang (School of Economics and Management, Beihang University, China)
Author(s): Shengjie Wang;Yanfei Kang

Combining multiple probabilistic forecasts has been known to improve forecasting accuracy. However, most existing combining research optimizes forecast metrics while neglecting their impact on downstream decision-making. In applications such as inventory management, decision-oriented measures—such as fill rate and inventory investment—are often more critical than conventional forecast accuracy metrics. Given the inherent trade-offs between forecasting accuracy and decision performance, optimizing a single metric can lead to overfitting and suboptimal decisions. To address this challenge, we propose a multi-objective probabilistic forecast combination framework that jointly optimizes both forecast accuracy and decision performance. Unlike traditional approaches focusing on a single objective, our method leverages multi-objective optimization to derive Pareto-optimal forecast combinations. We construct these combinations using linear pools and optimize model weights to obtain a diverse set of optimal solutions. A final selection is refined using the ideal point method, ensuring practical applicability in decision-making. Real-world empirical studies show that our method outperforms individual models and single-objective-based combinations by effectively balancing competing objectives. Furthermore, we analyze the Pareto-optimal combination weights, providing valuable insights into how our method reconciles conflicts between forecast accuracy and decision performance. These findings offer a robust framework for integrating probabilistic forecasts into real-world decision-making processes.

Micro-scale AI based Weather Prediction

Presenter: Jie Yan (North China Electric Power University)
Author(s): Jie Yan

Weather forecasting is the key input for wind power prediction. Currently, numerical weather prediction (NWP) based on solving system of partial differential equations is used in wind farms, which has problems of long computation time and low forecasting resolution. The AI-based weather forecasting can ensure prediction accuracy while improving the calculation speed. However, the existing AI-based weather forecasting models are mostly mesoscale weather forecasts with spatial resolution of more than tens of kilometers, which makes them impossible to accurately predict the wind speed at the location of wind turbines' hub. In order to solve the problems of low resolution and long calculation time in traditional numerical weather prediction, we proposed a microscale artificial intelligence wind speed prediction model (Micro-AWP) to serve wind power prediction. Vision Transformer(ViT) directly models the global information of an image, capturing long-range dependencies between different locations in the input image through a self-attention mechanism. Micro-AWP used ViT to build a small-scale fine-tuning network, and the ERA5 data was used to drive the medium-range weather forecasting. By leveraging macroscopic meteorological background information and microscopic local detail information of the wind farm, it realized the mapping from historical grid meteorological data to the future wind speed at the hub height of wind turbines, thus improving the accuracy of wind speed prediction within the wind farm.

Technical Trading: Is it Still Beating the Stock Market?

Presenter: Xinyue Hu (Zhongnan University of Economics and Law)
Author(s): Xinyue Hu;Yu Ren

In this paper, we further investigate the predictability and profitability of technical analysis in the stock market. We innovatively integrate Principal Component Analysis (PCA), Iterated Combination Constrained Predictor (ICCP) and Mallows Model Averaging (MMA) methods to construct a combined forecasting framework that addresses the limitations of traditional prediction models in factor selection. Based on this framework, we employ a large set of technical analysis indicators to conduct out-of-sample predictions for the weekly returns of major stock market indices in China. The empirical results show that while individual technical indicators have limited predictive power, the

selected combinations obtained through this method produce significantly better forecasting results than the benchmark model and outperform the commonly used naive equally weighted approach in traditional studies. Furthermore, the technical analysis-based investment strategy constructed from our predictions demonstrates significant excess returns, highlighting its practical economic value.

FinE2E-MTPOM: Financial End-to-End Multi-Task Portfolio Optimization Model

Presenter: Yupeng Wang (Capital University of Economics and Business)
Author(s): Yupeng Wang;Yinhong Yao;Zhensong Chen

In a highly dynamic and uncertain financial market, constructing an investment portfolio that can balance high returns with low risk has always been a core demand of investors and asset management institutions. Recently, deep learning-based methods have gained increasing attention in addressing this challenge. However, these methods usually handle stock prediction and investment decision-making in separate stages, leading to a mismatch between the optimization target and the actual demand. This paper proposes a financial end-to-end multi-task portfolio optimization model, referred to as FinE2E-MTPOM, which significantly improves the return-risk balance of the investment portfolio by integrating stock prediction and investment decision-making. Specifically, it adopts an end-to-end training mechanism to achieve a two-way interaction between stock prediction and investment decision-making, and directly optimizes the model parameters according to the investment target. The core technology of the model includes four parts: 1) Extracting stock features through time series learning components to capture the short-term and long-term time series characteristics of stocks; 2) Applying the Self-supervised Graph Attention Network to learn the implicit relationship between individual stocks and sectors; 3) Introducing a multi-task prediction framework to simultaneously predict stock volatility risk, return ratio, and movement; 4) The stock investment proportion is optimized using a minimax problem, which is reformulated through convex duality theory for efficient solution within the end-to-end system. Experiments on the S&P 500 and CSI 300 data sets demonstrate that the proposed FinE2E-MTPOM model outperforms state-of-the-art methods in key evaluation metrics, including cumulative return, Sharpe ratio, and maximum drawdown compared to state-of-the-art methods, verifying its significant advantages in real-world investment scenarios.

Research on Generalized Leverage Effect in Market Volatility

Presenter: Jiafu Xu (University of International Business and Economics)
Author(s): Jiafu Xu;Haibin Xie

Although the leverage effect in volatility has been widely studied, existing literature almost implicitly assumes that the leverage effect is not affected by the level of volatility. This paper designs a model structure with a generalized leverage effect to verify this implicit assumption. Based on the transaction price data of different countries' stock markets, we find that there is a very significant generalized leverage effect in the stock market, that is, the leverage effect is significantly affected by the market's own volatility level. The empirical results in and out of the sample show that the volatility model considering the generalized leverage effect can significantly improve the prediction accuracy of market volatility. The research findings of this paper not only provide a new perspective for understanding the leverage effect, but also provide new empirical facts for volatility modeling.

Optimizing Energy Futures Investments under Uncertainty Shocks: a Multi-input Multi-output Forecasting Framework based on Multi-objective Optimization and Deep Learning

Presenter: Weixin Sun (Dongbei University of Finance and Economics)
Author(s): Weixin Sun;Yong Wang;Li Zhang;Xihui Haviour Chen

Energy price forecasting is vital for ensuring energy security and optimizing resource allocation. However, recent uncertainty shocks, such as geopolitical tensions, climate change, and public health crises, have intensified price volatility, increasing forecasting complexity. This study proposes an enhanced mixed-frequency data fusion multi-input multi-output deep learning network (EMF-MIMO-DLN) to improve forecast accuracy and reliability. The framework integrates four modules: mixed-frequency data fusion, feature extraction, forecasting, and performance

evaluation. Specifically, the reverse restricted mixed data sampling (RRMIDAS) method is used to incorporate low-frequency uncertainty indices into high-frequency energy price prediction, enhancing responsiveness to economic fluctuations and shocks. The integrated indices specifically include: climate policy uncertainty (CPU), the global geopolitical risk index (GPR) with country-specific variants for Russia (GPR_RUS), the U.S. (GPR_USA), and Saudi Arabia (GPR_SAU), alongside monetary policy uncertainty (MPU), oil price uncertainty (OPU), trade policy uncertainty (TPU), and the US-China tension index (UCT). Building on the advanced patch time series Transformer (PatchTST) model, the framework is designed with a multi-input multi-output (MIMO) structure that provides both point and interval predictions, strengthening risk management capabilities. Additionally, this study introduces an intelligent interval constraint investment strategy, which effectively improves the reliability of investment strategies and mitigates the investment risks arising from forecasting uncertainty. Through a series of comparative experiments and validations conducted on the crude oil and natural gas markets, the study demonstrates the significant advantages of the proposed forecasting framework in improving prediction accuracy, addressing uncertainty shocks, and optimizing investment decisions, showcasing its broad application potential in the energy market.

Boilover of Crude Oil Storage Tanks: Mechanism, Prediction, and Suppression

Presenter: Yuntao LI (China University of Petroleum, Beijing)
Author(s): Yuntao Li

Boilover ranks as one of the most catastrophic scenarios in tank fires, characterized by intense combustion and substantial oil spillage. The heat release rate and flame height increase dramatically, posing a substantial threat. Water is essential for boilover to occur, with both the inherent water content in crude oil and the water layer at the bottom of the tank acting as triggers. Boilover is highly sensitive to various factors and boundary conditions, making it difficult to prevent and control. This study focuses on the key factors influencing boilover fires, examining the impact mechanism and uncertainty analysis of crude oil water content, the influence of tank ullage height and water layer thickness on boilover behavior, and the role of tank boundary conditions in shaping boilover characteristics.To complement this research, large-scale experiments with a diameter of 1.5 meters were specifically conducted, enriching the theoretical understanding and prediction of boilover. The study also explores boilover suppression technologies, including experiments using self-developed solid spheres for oil surface coverage and suppression, as well as combined fire-extinguishing tests with solid spheres and foam. The findings provide valuable insights for combating full-surface fires in crude oil tanks.

Measuring and Nowcasting China′ s Growth-at-Risk with High-Dimensional Stock Market Connectedness

Presenter: Kunchao Li (Shanghai University of Finance and Economics)
Author(s): Kunchao Li;Shuting Liu;Shiqi Ye

As a barometer of macroeconomic conditions, the stock market contains a vast amount of information that provides a critical channel for characterizing economic dynamics and offers a potential avenue for monitoring and early warning of macroeconomic risks. This study integrate high-dimensional stock market connectedness, mixed-frequency machine-learning techniques, and the Growth-at-Risk (GaR) framework into a unified empirical setting to measure and monitor China's growth-at-risk in real time. Specifically, we first characterize the highdimensional dynamic correlation structure among 1250 stocks using the VDcDCCGARCH model. Based on this, firmlevel topology indicators are extracted from the minimumspanningtree (MST) network. These indicators are then incorporated into the group-regularized mixed-frequency quantile regression model, enabling the dynamic measurement and real-time warnings of China's GaR. Empirical findings reveal that incorporating high-dimensional stock market connectedness into the GaR framework not only promptly captures growth vulnerabilities—particularly in response to major shocks such as the COVID-19 outbreak—but also significantly enhances the accuracy of real-time forecasts compared to conventional approaches. Moreover, the contribution of stock-level information to growth risk exhibits pronounced heterogeneity across sectors. Specifically, stock-level information of the consumer staples and consumer discretionary sectors emerge as key predictors of China's macroeconomic downside risk. Against the backdrop of a rapidly evolving global landscape and the confluence of extreme events, this study provides a novel toolkit for the identification and early warning of macroeconomic vulnerabilities. It also offers important empirical evidence and policy guidance for

safeguarding economic stability.

Fire Risk Level Prediction for Ancient Building Complex of the Palace Museum in China Based on Machine Learning Models

Presenter: Zhanfeng Shen (Aerospace Information Research Institute, Chinese Academy of Sciences)
Author(s): Zhanfeng Shen;Xinwei Yang;Yating Lei

The complexes of historic timber-structure buildings are extremely complicated and have various fire threats. Therefore, how to prevent fires from occurring is one of the most critical issues faced in the fire prevention and safety protection of ancient buildings. In this study, we firstly build a fire risk assessment indicator system that emphasizes the specificities of timber and historic structures. And then a fire risk index-driven solution is proposed based on machine learning models (namely are XGBoost, GBDT, CatBoost, Adaboost, and SVM) to predict the fire risk level of timber heritage buildings located in the Palace Museum in China. Through the proposed solution, the monthly fire risk level of the ancient buildings of the Palace Museum in 2019 was assessed by machine learning algorithms using the fire risk index data of the historical years (2016–2018). The accuracy evaluation showed that XGBoost performed better than other models, which is expected to provide information support for the safety inspection work of the fire department in the Palace Museum.

Price Level and Financial Crisis Prediction

Presenter: Li Zhengyang (Fudan University)
Author(s): Zhengyang Li;Zaichao Du

Financial crises, among the most disruptive economic events, necessitate precise identification and forecasting of their long-term output losses and systemic risk transmission mechanisms. Traditional approaches, constrained by limited crisis samples, lagging indicators, and the “black-box” nature of machine learning (ML), struggle to achieve causal inference and policy applicability. This study pioneers the application of the Double Machine Learning (DML) framework to financial crisis prediction. By integrating 236 macroeconomic variables across eight categories (e.g., composite indices, credit markets, political environment) from 17 OECD countries (1950–2016) through heterogeneous databases (JST, WIOD, etc.), we construct a multi-source dataset and innovatively adopt recursive window training with localized normalization to mitigate look-ahead bias.Key findings include: (1) Household and government consumption price levels emerge as robust predictors of financial crises, with significant nonlinear relationships to crisis severity based on historical percentiles; (2) The DML framework, leveraging Neyman-orthogonal scores and cross-fitting, successfully disentangles complex, endogenous relationships, revealing causal pathways related to credit expansion and asset price bubbles; (3) While dimensionality expansion using Random Fourier Features (RFF) does not always improve out-of-sample performance, the Nyström method proves effective in mitigating overfitting in overparameterized regimes, demonstrating its value in enhancing model generalization. Empirical results suggest that incorporating novel Early Warning Indicators (EWIs) based on consumption price levels enables more accurate and timely crisis predictions, supporting proactive policy interventions. This research advances the causal interpretability of ML in financial stability, highlighting the importance of integrating new variables that capture the dynamics of economic instability. It also establishes a robust methodological framework for high-dimensional factor models in economic forecasting, offering a more reliable foundation for policy decisions and improving the precision of financial stability assessments across diverse economic environments. The identification of new predictors contributes significantly to the literature, providing deeper insights into the dynamics of financial crises and their early identification.

Multi-modal Fusion Based on Large Language Model and Ensemble Learning - A Case Study in SME Financing Guarantee Evaluation

Presenter: Cui Fenghao (Beijing Technology and Business School)
Author(s): Fenghao Cui;Jiaming Liu

This project aligns with China's "Inclusive Finance" strategy, focusing on the intelligent enhancement of financing guarantee decision-making systems for small, medium, and micro-enterprises (SMMEs). As vital pillars of the national economy, SMMEs drive market vitality, foster technological innovation, and expand employment opportunities. However, fragmented information disclosure limits available numerical features, while traditional "black-box" models obscure risk transmission pathways, undermining the credibility of assessments in practical scenarios. To address this, we develop a LIME-based data augmentation algorithm. Unlike existing methods reliant on simplistic concatenation, which struggle with shallow interaction bottlenecks and fail to capture deep semantic links across financial data, enterprise texts, and industry policy documents, we propose an improved multi-classifier ensemble prediction method using dual-modal information. Additionally, given SMMEs' risk evaluation heavily depends on domain expertise—due to their small scale, volatile operations, and elusive risk signals in data—we introduce an LLM Agent-based model integrating local and historical experiential insights for financing guarantee assessments.

Sentiment-driven Insights: Unlocking the Sentiment Dynamics and Forecasting the Future of Crude Oil Prices with Big Language Models

Presenter: Mingchen Li (China University of Mining and Technology)
Author(s): Mingchen Li;Yunjie Wei;Yuanyuan Shi;Shouyang Wang

Accurate crude oil price forecasting is essential for effective energy policy and market strategies. This study presents an innovative framework that integrates structured time-series data with unstructured textual inputs, such as news articles, financial reports, and social media sentiment, to enhance forecasting accuracy. A state-of-the-art large language model (LLM) is employed to extract nuanced sentiment indicators from unstructured data, marking a novel application of advanced natural language processing in this field. The framework's core innovation lies in leveraging Temporal Convolutional Networks (TCN), a deep learning architecture optimized for capturing long-term temporal dependencies in time-series data. TCN's unique design addresses limitations of traditional models and alternative neural networks, offering superior performance in handling sequential data. This hybrid approach effectively combines diverse data sources and cutting-edge predictive techniques, showcasing the potential of integrating structured and unstructured data for complex forecasting tasks.

Can Large Language Models Evaluate Suppliers′ qualifications of Complex Production Lines? An Approach Incorporating Retrieval-augmented Generation and Chain-of-Thought

Presenter: Xuerong Li (Academy of Mathematics and Systems Sciences, Chinese Academy of Sciences)
Author(s): Xuerong Li

In recent years, large language model (LLM) technology has introduced new opportunities to enhance the efficiency and accuracy of supplier qualification and risk management. Through experiments utilizing a unique dataset from a large Chinese defense enterprise's supplier base, this study finds that LLMs can effectively perform hierarchical evaluation tasks based on multidimensional indicators. Compared to traditional indicator-based evaluation models, LLMs not only provide relatively accurate evaluation levels but also generate auxiliary management comments rapidly and automatically. Additionally, by integrating RAG and COT techniques into the LLM framework, the professional domain knowledge and evaluation capabilities in specific areas can be further enhanced, thereby improving the accuracy of hierarchical evaluations.

Assessment of Climate Policy Uncertainty Using Large Language Models and Energy Market Interval Forecasting

Presenter: Liu Jiajia (Beijing Univeristy of Chemical Technology)
Author(s): Jiajia Liu;Weiyi Zhang

This study quantifies the uncertainty of global climate policies through large language models and delves into their impact on the energy market, with a focus on oil prices. By integrating CEEMDAN and BNN-STGCN technologies, we've developed a multi-scale analysis framework. It can handle complex nonlinear time series data

and capture the energy market's spatio-temporal dependencies. Results indicate that climate policy uncertainty is a significant predictor, especially during times of frequent policy changes or high market volatility. Furthermore, the use of CEEMDAN and BNN-STGCN technologies has greatly enhanced the precision of short-term and medium-term forecasts. This offers valuable insights to market players and policymakers about how policy shifts influence the energy market. Finally, Bayesian estimation has been incorporated to realize the interval prediction of crude oil prices. The methods and outcomes of this study provide a scientific basis for risk management and strategy formulation in the energy market. They also highlight the potential uses of large language models and Bayesian estimation in macroeconomic analysis.

Forecasting the Yield Curve with the Time-Varying Neural Dynamic Nelson-Siegel Model

Presenter: Sicco Kooiker (Vrije Universiteit Amsterdam)
Author(s): Sicco Kooiker;Julia Schaumburg;Janekke Brummelen

In this article, we introduce neural network-predicted factor loadings in the dynamic Nelson-Siegel yield curve model for simultaneous analysis and forecasting of interest rates across different maturities. The classical Nelson-Siegel model has been extended to include vector autoregressive factors and time-varying factor loadings, yet the functional form of these loadings has remained arbitrary despite their recognized importance. We address this limitation by developing a flexible approach that learns factor loadings through two shallow neural networks. Our simulation study shows that a nine-parameter neural network can closely approximate traditional Nelson-Siegel factor loadings. These neural networks improve both in-sample and out-of-sample filtered estimates of the US Treasury bond yield curve. Recognizing the similarity between Stochastic Gradient Descent and score-driven filters, we extend our approach by allowing neural network parameters to evolve in a score-driven manner. We integrate these dynamics within the Kalman filter state space framework to create a simultaneous optimization routine. Our empirical study demonstrates the validity of our method compared to traditional approaches in yield curve forecasting up to one year ahead.

Shadow Banking, Threat or Opportunity for the Developing Countries

Presenter: Mariam Lashkhi (East European University)
Author(s): Mariam Lashkhi

Shadow banking, a complex web of non-bank financial intermediaries, presents a multifaceted challenge for developing countries. This sector, operating outside traditional regulatory frameworks, offers potential benefits like increased financial inclusion and alternative funding sources, particularly crucial in economies with underdeveloped formal banking systems. However, it also introduces significant systemic risks, including regulatory arbitrage, maturity mismatches, and a lack of transparency, potentially exacerbating financial instability. Developing nations often grapple with limited access to credit and capital, hindering economic growth. Shadow banking can bridge this gap by providing innovative financial products and services, catering to underserved populations and small and medium-sized enterprises (SMEs). This sector's agility and flexibility can foster entrepreneurship and stimulate local economies. Furthermore, it can attract foreign investment and diversify funding sources, reducing reliance on traditional bank lending. However, the opaque nature of shadow banking poses substantial risks. The absence of stringent regulatory oversight can lead to excessive leverage, asset bubbles, and contagion effects, replicating the vulnerabilities that triggered the 2008 global financial crisis. Developing countries, with weaker regulatory capacities, are particularly susceptible to these risks. The lack of deposit insurance and lender-of-last-resort support can amplify the impact of financial shocks, potentially triggering widespread financial instability. Furthermore, shadow banking can facilitate illicit financial flows, including money laundering and tax evasion, undermining governance and hindering sustainable development. The potential for regulatory arbitrage, where entities exploit regulatory loopholes, can erode the effectiveness of national financial regulations. To mitigate these risks and harness the potential benefits, developing countries must adopt a balanced approach. This involves strengthening regulatory frameworks to enhance transparency and oversight, while promoting financial inclusion and innovation. International cooperation is crucial to address cross-border risks and establish global standards. By carefully navigating the complexities of shadow banking, developing countries can leverage its potential to drive economic growth while safeguarding financial stability.

Forecasting the Term Structure of Government Bond Yield Using B-Splines with Time-varying Control Points and Knots Positions

Presenter: Xia Zou (Vrije universiteit Amsterdam and Tinbergen Institute)
Author(s): Xia Zou;André Lucas

This paper introduces a Flexible Dynamic B-Spline (FDBS) yield curve model, which extends the dynamic B-spline approach (DBS). In contrast to the DBS model, which allows only control points to vary over time, the FDBS model allows both control points and knot positions to evolve over time. Similar to the Dynamic Nelson-Siegel (DNS) model by Diebold and Li, the FDBS model imposes a factor structure on the yield curve, where factor loadings are specified as B-spline basis functions and factors correspond to the control points. A key feature of the FDBS framework is that both factors and factor loadings are time-varying, with the dynamics of factor loadings efficiently captured through a small number of inner knot positions. To maintain the strict ordering of knots, we apply a logistic normalization transformation. The joint evolution of control points and knot positions is modeled within the Generalized Autoregressive Score (GAS) framework. Through a simulation study, we demonstrate that under a data-generating process (DGP) where only control points vary over time, the FDBS model performs comparably to the DBS model. However, when both control points and knot positions are time-varying, the FDBS model significantly outperforms the DBS model. Empirical analysis using U.S. Treasury zero-coupon yields from January 1985 to December 2009 (with 17 fixed maturities) further confirms the advantages of the FDBS model. We show that the FDBS models with one, two, and three inner knots outperform both the DBS model with same amount of inner knots and the DNS model—in terms of both in-sample and out-of-sample performance. And the results show most of the improvements come from forecasts at long horizons. Additionally, out-of-sample results indicate that the FDBS model with two inner knots provides the best forecasting accuracy.

Measuring Risk Perception for Copper Commodity Price Forecasting

Presenter: Xuefei Wu (Academy of Mathematics and Systems Science, Chinese Academy of Sciences)
Author(s): Xuefei Wu;Wei Shang;Lei Cao

Nonferrous metals, as critical natural resources for national economies and industrial production, require timely market monitoring to inform production decisions and optimize pricing mechanisms for energy financial derivatives. Addressing the challenge of forecasting copper futures prices, this study proposes an innovative hybrid modeling framework that integrates multi-dimensional market sentiment and risk factors, aiming to establish a scientific analytical paradigm for commodity price prediction.Using text mining methods, this research extracts market-oriented sentiment signals from industry portals to decode the laws governing supply-demand dynamics and price movements in nonferrous metals markets. A comprehensive domain-specific lexicon framework is constructed to classify and evaluate thematic terms, trend indicators, and emotional lexemes, assigning differentiated weights based on their contextual relevance and market impact intensity. Rooted in the BERTopic classification system, the designed market sentiment index system transcends the “supply-demand-price” three-dimensional paradigm, establishing a four-dimensional framework covering supply, demand, financial environment, and policy orientation. For risk variable identification, a multi-modal feature extraction module employs a dual-track approach: explicit risks (e.g., policy adjustments, geopolitical conflicts) are captured via keyword matching and regular expression techniques, while implicit risk signals in unstructured text are mined using pre-trained semantic models to calculate cosine similarity with a pre-built risk lexicon. This mechanism enhances the comprehensiveness and proactivity of risk monitoring, integrating multi-faceted risk information into the predictive model.Empirical results show that the four-dimensional sentiment classification system outperforms baseline models in forecasting accuracy. The hybrid model combining risk and sentiment variables achieves significantly more precise predictions of copper futures closing prices than benchmark models. During periods of intense market risk volatility, risk perception indicators demonstrate superior predictive efficacy to traditional sentiment variables, enabling more sensitive capture of early warning signals for shifts in market supply-demand structures and price fluctuations.

Insights into Cross-category Effects in Demand Forecasting

Presenter: Ozden Gur Ali (Koc University)
Author(s): Ozden Gur Ali

Estimating the impact of promotions and forecasting demand in grocery stores is important to retailers for both promotion planning and procurement decisions. This task is challenging due to multiple issues: Potential confounding due to factors including seasonality; the short- and longer-term effect of promotions on customer behavior – such as increasing demand during the promotion while potentially decreasing it before and after the campaign; as well as cross-product and cross-category effects on demand. Further, promotions are not conducted as random designed experiments but coordinated across categories and products to maximize their impact.

FairRSF: Fairness-aware Retail Sales Forecasting Via Prototype Decomposition

Presenter: Li Zhou (Hefei University of Technology)
Author(s): Li Zhou

Sales forecasting plays a crucial role in retail management. While existing studies focus on improving forecasting accuracy, fairness issue is largely under-investigated, leading to accuracy disparities across variables, such as products and retail stores. Unfair retail forecasting may induce inequitable resource allocation and strategic disadvantages for marginalized variables. To achieve fair retail sales forecasting, we present a novel end-to-end deep learning method named FairRSF. The FairRSF method is mainly composed of two components, i.e., (1) multi-level sales pattern decomposition to uncover fairness-sensitive factors through discrete wavelet transform and prototypical contrastive learning, and (2) fairness enforcement mechanism integrating adversarial learning and a fairness-aware graph neural network. Comprehensive experiments on real-world retail datasets demonstrate that FairRSF achieves comparable forecasting accuracy and significantly improves forecasting fairness across variables.

Advancements in the Application of Interpretable Models for Retail Purchase Prediction

Presenter: Liu Zhenkun (Nanjing University of Posts and Telecommunications)
Author(s): Zhenkun Liu

Current research in retail forecasting primarily focuses on the macro level, with an emphasis on sales forecasting. In contrast, this presentation shifts the focus to the micro level, specifically predicting individual customer purchase behavior. By forecasting a customer's willingness to purchase a product, retailers can achieve more precise personalized marketing. With the introduction of interpretable predictive techniques in customer purchase intent forecasting, retailers can gain deeper insights into the underlying factors driving individual purchasing decisions, allowing for the development of more tailored marketing strategies for different customer segments. This presentation will highlight the progress of interpretable models in retail purchase prediction, categorizing their evolution into three stages: (1) simple models with transparent structures, (2) classic ensemble models with complex structures, and (3) customized ensemble models with complex structures. By tracing the development of retail purchase prediction from simple models to complex ensemble approaches, this presentation aims to demonstrate how researchers have leveraged innovative technologies to progressively improve prediction accuracy while maintaining the interpretability of the models.

Retail Store-SKU level Replenishment Planning with Attribute-Space Graph Recurrent Neural Networks

Presenter: Shaohui Ma (Nanjing Audit University)
Author(s): Shaohui Ma

Large retail chains routinely make replenishment decisions for inventory at the store Stock-Keeping-Unit (SKU) level, relying heavily on accurate demand forecasting to meet target service levels. We critically examine the current state-of-the-art in demand forecasting, revealing a notable gap in existing models, specifically the neglect of intricate

cross-SKU promotional effects. To address this gap, we introduce an innovative solution: the Attributed Space Graph Recurrent Neural Network (ASG-RNN), which encodes demand history and drivers' history using a recurrent neural network, and captures cross-SKU promotional interactions through a graph convolution network. The graph convolutions operate on an attribute space graph constructed using K-nearest neighbor algorithm based on SKU-level static attributes. Initial empirical tests were conducted on a dataset comprising 200 SKUs in the Coffee category from a retail chain with 93 stores. The analysis demonstrates that ASG-RNN, on average, reduces loss by 4.1% compared to the best benchmark, resulting in average inventory cost savings of 4.94%. Notably, for promoted SKUs, ASG-RNN reduces loss by 13% on average, and it can achieve a 30% reduction in loss when the target service rate is high. Further experiments conducted on four additional product categories, including Beer, Laundry Detergent, Milk, and Mayo, yield similar results, showcasing the robustness of ASG-RNN's performance. These findings underscore the practical significance and cost-effectiveness of our innovative approach in enhancing demand forecasting and optimizing replenishment decisions for large retail chains.

A Report of Forecasting Visitor Arrivals after COVID-19 Pandemic

Presenter: Jialu Gao (Macau University of Science and Technology)
Author(s): Jialu Gao;Linyue Zhang;Jiyang Wang;JIngrui Li;Shoujiang Li

Given the impact of the COVID-19 pandemic on the global tourism industry, it is important to predict the number of inbound and outbound tourists during the recovery period of the epidemic. In this study, econometric and artificial intelligence methods are combined to forecast China's tourist arrivals in the coming year. The number of tourists in 20 regions is processed by point forecast and interval forecast, and the forecast results are adjusted accordingly with reference to the actual situation under different recovery scenarios. The experiments prove that, compared with the single models, the combined model is more helpful to improve the accuracy and robustness of tourism demand forecasting, and help policy makers make reasonable tourism decisions.

Tourism Forecasting Competition Amid COVID-19 Round II: Recovery of the Chinese Outbound Travel Market from the Pandemic

Presenter: Baobao Song (University of Technology Sydney)
Author(s): Baobao Song;Yishuo Zhang;Gang Li

Tourism demand forecasting remains a critical challenge in the tourism industry, especially in the aftermath of external disruptions such as the Covid-19 pandemic. Accurate forecasts are essential for policymakers and businesses to develop effective recovery strategies. Traditional econometric and statistical models, such as ARIMA, depend on stable historical trends. However, the Covid-19 pandemic caused significant disruptions, fluctuations, and misalignments in tourism data, rendering pre-pandemic patterns unreliable and requiring modifications to existing forecasting methods. To address these challenges, we propose a novel approach that incorporates tourism data imputation and a two-stage forecasting framework. The data imputation step helps resolve inconsistencies across various factors, ensuring better alignment and continuity in the dataset. In the first stage, historical tourism data from the pre-pandemic period is used to forecast demand during the pandemic and the early recovery phase. In the second stage, the focus shifts to the pandemic period to estimate the recovery rate, which is then used to adjust the demand forecasts for the post-pandemic period. This approach not only enables more robust predictions of future tourism demand but also provides valuable insights into the recovery process. As such, it offers a practical and reliable tool for tourism forecasting in the post-pandemic era. This research contributes to the existing literature by introducing a comprehensive framework that helps tourism stakeholders navigate the uncertainties of recovery more effectively.

Destination-Specific Interpretable Learner (DSIL): A Novel Approach for Efficient and Interpretable Tourism Demand Forecasting in the Post-COVID Era

Presenter: Zhang Huihui (The Hong Kong Polytechnic University)
Author(s): Huihui Zhang;Hao Li;Florian Zach;Zheng Xiang

The COVID-19 pandemic has profoundly disrupted global tourism, making accurate forecasting of tourism

demand more critical yet more challenging than ever. This report introduces a novel forecasting model, the Destination-Specific Interpretable Learner (DSIL), designed to produce accurate, interpretable, and efficient tourism demand predictions in the post-pandemic era. Addressing limitations of traditional linear models and opaque deep learning approaches, DSIL combines Seasonal-Trend decomposition using Loess (STL) and Light Gradient Boosting Machine (LightGBM) to capture seasonality and complex, nonlinear relationships within the data. Leveraging tourism arrival data and a wide array of covariates—including exchange rates, pandemic indicators, Baidu search trends, and economic indices—the model customizes forecasts for 20 destinations, each with distinct tourism drivers. The study introduces a quantile regression framework to forecast not just point estimates, but also prediction intervals, offering stakeholders a clearer understanding of future demand uncertainties. Time-series cross-validation, using rolling origin splits, ensures robust model validation across both pre-pandemic and pandemic-affected periods. DSIL achieves strong forecasting performance, with Mean Absolute Percentage Error (MAPE) values between 10% and 30%, outperforming expectations given the volatility of long-range tourism forecasting. Results are particularly accurate for destinations heavily visited by Chinese tourists, such as Hong Kong, Macau, and Thailand. The model also identifies destination-specific variable importance, offering valuable strategic insights. For instance, features like exchange rate and COVID-19 vaccination rates were highly influential in predicting Chinese tourist flows to the U.S., while others like unemployment were less impactful in this context. The report concludes that DSIL offers a significant advancement in tourism forecasting by balancing interpretability, computational efficiency, and predictive accuracy. While current limitations include a relatively narrow set of variables and lack of direct benchmarking against baseline models like ARIMA or SVM, future extensions plan to address these. Overall, DSIL holds strong practical relevance for policymakers and industry stakeholders seeking to navigate the evolving landscape of global tourism recovery with data-informed strategies.

Resilience Assessment of the Energy Industry Chain from the Production and Consumption Perspectives

Presenter: Shuning Gao (Xidian University)
Author(s): Shuning Gao;Jian Chai

Against the backdrop of intertwined global energy transition and geopolitical risks, the resilience of the energy industry chain has become a core issue in safeguarding national energy security. The study selects the panel data of 30 provinces in China from 2005 to 2022, constructs a resilience evaluation index system of the energy industry chain covering “robustness, resilience, and adaptability”, and comprehensively applies spatial econometric modeling and dynamic cluster analysis to reveal the resilience patterns and driving mechanisms of the energy-producing and energy-consuming provinces. To determine how resilient the provinces that produce and consume energy are, the system will be examined using a geographic econometric model and dynamic clustering analysis. The study found that: (1) The energy-production resilience regions are dominated by resource endowments, with the resource-rich western provinces taking advantage of the production resilience by the energy self-sufficiency rate and the diversification index, but the over-reliance on inter-provincial energy transfers leads to an increase in the vulnerability of the industry chain. (2) The energy-consumption resilience regions show a differentiated pattern of market-regulated demand, with eastern provinces realizing consumption resilience by reducing energy intensity and investment in innovation, forming a transition path of “high adaptation - low robustness”; while the industrial-intensive regions in central and western China are constrained by the negative spillover effects of industrial agglomeration, and the potential for energy efficiency improvement is limited. (3) Regional synergies are significant, but the efficiency of synergies varies considerably among “production-consumption” provinces. Scenario analysis suggests that the implementation of the differentiated strategy of strengthening storage and transportation interconnections in the west and deepening market reforms in the east could narrow the regional resilience gap. The study develops a cross-scale analytical framework from the standpoint of “production-consumption” synergy, revealing the deeper mechanism of regional resilience imbalance, which is critical to the global energy transition and regional policy formulation under China’s dual-carbon strategy.

The Impact and Prediction of a News Media-based Climate Risk Attitudes on China's New Energy Stock Market

Presenter: Limin Xing (Xidian Univesity)
Author(s): Li-Min Xing;Qi-Wen Xiao;Jian Chai

Global climate change has become a significant challenge affecting global economic and social development, with the renewable energy industry being strongly promoted by policies worldwide. However, the renewable energy stock market exhibits high volatility under the influence of multiple factors. This study constructs a Climate Risk Attitude (CRA) index based on multi-source integrated news media data, using the BERT deep learning model and SnowNLP sentiment analysis tool. By combining ARDL, NARDL, and QARDL models, the study analyzes the linear, nonlinear, and asymmetric effects of climate risk attitude on the renewable energy stock market. Additionally, an LSTM model is employed to evaluate the predictive effect of climate risk attitude on the renewable energy stock market. The results show that climate risk attitude has a significant impact on the renewable energy stock market, and its inclusion can effectively improve the prediction accuracy of the market, providing important theoretical and empirical support for renewable energy investment decisions.

Resilience Measurement and Risk Prediction in Critical Mineral Supply Chains: Leveraging Renormalization Theory and Geometric Bifurcation Growth Models

Presenter: Qian Zhi (Xidian University)
Author(s): Qian Zhi

The resilience of critical mineral supply chains is a central issue in international competition and the green transition. This study constructs a dynamic multi-layer complex network based on 2,387 global supply chain datasets of new energy and critical mineral enterprises from January 2017 to December 2024. Leveraging renormalization theory and geometric bifurcation growth models, we quantify the resilience of critical mineral supply chains, predict risk pathways under geopolitical interventions, and propose multi-scale optimization strategies. The research spans three hierarchical layers of mineral supply chains: raw materials, processed products, and final products. Nodes include enterprises and national entities from China, the United States, Australia, Japan, and other key countries. Through time-windowed analyses, we identify policy-induced lag effects on corporate partnerships, with varying response speeds across different supply chain layers. Using percolation theory, we evaluate critical nodes and assess network robustness, revealing vulnerabilities in key resource-dependent nations. For instance, removing critical nodes such as Australia (lithium) drastically diminishes the largest connected subgraph, highlighting risks of single-point dependencies. To decode cross-scale resilience mechanisms, we employ a renormalization model for multi-scale analysis. Using box-counting methods, the network is coarse-grained into national alliance and enterprise layers, confirming its self-similar fractal structure. Geopolitical simulations further demonstrate that removing U.S. ally nodes triggers a significant decline in global supply chain connectivity, indicating that geopolitical alliances amplify cascading risks. Through an inverse-renormalization geometric bifurcation growth model, we simulate supply chain evolution under U.S. policy interventions. Bidirectional scale transformations validate the self-similarity of supply chain networks, and scenario testing predicts risk propagation pathways. Results show that policy shocks significantly increase the probability of cascading failures in Chinese processing enterprises. While supply chain restructuring can mitigate risks, it incurs additional costs, such as rerouting mineral flows through alternative regions. Key findings urge nations to build multi-polar mineral alliances and firms to diversify suppliers, reducing cascading risks.

Making Black-box Forecasts Understandable: The Effect of Global and Local Explanations

Presenter: Kai Hoberg (Kühne Logistics University gGmbH)
Author(s): Kai Hoberg;Naghmeh Khosrowabadi;Jan Fransoo

Advanced Artificial Intelligent (AI) systems are becoming ubiquitous in demand planning and forecasting. While these AI systems are capable of processing massive amounts of data to generate more accurate forecasts, their lack of transparency can hinder human planners to confidently follow the generated recommendations. Distrust in system forecasts adversely affect the human-machine interaction, as it increases the chance of making unnecessary, inaccurate,

and detrimental adjustments. Explainable AI (XAI) aims to enhance the user's understanding of AI recommendations. Incorporating XAI existing techniques with an online sales forecasting task, we examine two types of explanations, i.e., global explanation and local explanation. While global explanations yield overall insights into the effects of different features on forecasts, local explanations clarify the effects in more detail for a given forecast. We test these two explanations in a laboratory setting to understand the impact of subjects' confidence in the forecasts provided. We find that providing explanations can preserve subjects' confidence in their decisions. In particular, we find evidence for a better performance of the global explanation, as subjects provided with this explanation selected not only shorter ranges, indicative of higher confidence in the estimates, but also less deviated ranges from the AI forecasts. While the global explanation successfully enhanced confidence in AI generated forecasts, the local explanation adversely added to overoptimism bias.

Does Frequency Bias in Lay Inflation Judgments Depend on Monitored Prices or Number of Purchases?

Presenter: Xiaoxiao Niu (Shenzhen University)
Author(s): Xiaoxiao Niu;Nigel Harvey

Previous research has identified a frequency bias in lay people's inflation judgments. This means that consumers' perceptions of inflation rates are systematically biased toward the perceived price changes of the frequently purchased items. However, limited attention has been paid to the effects of prices that consumers merely monitor rather than actually purchase. This study investigates whether it is the number of times that people are exposed to different prices that matters. In other words, does a frequency bias occur even when prices are merely monitored? If it does, people's expectations of inflation would be strongly influenced by a categories of item (e.g., new vehicles) that they purchase very rarely. In our experiment, we adapted the experimental design of Georganas et al. (2014). Four hundred and forty-five participants were recruited via the web platform Prolific. In a simulation of actual shopping experiences over a period of six months (96 days), participants were randomly allocated to either a purchase condition or a monitor condition. Afterwards, they made inflation judgments for each of the four products and for the basket over the six months. Our results replicated Georganas et al. (2014)'s finding of frequency bias and showed that people put, on average, 56% weight on the frequency with which goods are purchased or monitored and 44% weight on the (theoretically correct) expenditure weights. Additionaly, frequency bias existed in both the purchase condition (75%) and the two monitor conditions (monitor the high-price good with relatively low frequency—41%, monitor the low-price good with relatively higher frequency—56%). Results also showed that people's judgment errors for good-specific inflation were significantly affected by inflation level (when the basket inflation is the same, whether the good-specific inflation rate is positively/negatively/not related to the frequency of seeing the specific good), but not influenced by different tasks (purchase versus monitor conditions). Our findings confirmed the frequency bias in laypeople's inflation judgment and demonstrated that frequency bias depends both on the number of purchased prices as well as the number of monitored prices. To better forecast lay people's inflation judgment, relevant departments and organisations should not only collect expenditure information but also make an effort to access those categories/items people usually monitor but rarely buy.

Judgmental Model Pooling

Presenter: Fotios Petropoulos (University of Bath)
Author(s): Fotios Petropoulos;Evangelos Spiliotis;Ville Satopaa

Time series forecasting often incorporates human judgment in the form of model selection, model parametrization or forecast adjusting. In this paper we explore for the first time the potential of judgment in model pooling, a process that involves selecting one or multiple appropriate models from a pool of available models to improve accuracy. Through a behavioural experiment, we assess how human judgment performs in model pooling compared to a statistical model selection and a simple model combination approach, and identify factors influencing this performance. Our results indicate that judgmental model pooling surpasses both benchmarks when two or more models are combined and the forecasts from all possible forecast combinations are visualised to the forecasters. Similarly, judgmental model pooling tends to outperform statistical model selection in series of weak trend/seasonality, doing also better than the simple combination in series of strong trend/seasonality. Another interesting finding is that humans

are more effective in rejecting the worst models from their pools than the statistical selection, thus improving the robustness of their forecasts. In addition, we find evidence that combining more models, getting familiar with the task, and receiving feedback about the performance achieved have positive effects on judgmental model pooling. We discuss the implications of our findings for forecasting practice and list possible limitations and advances.

High-Dimensional Conditional Factor Model

Presenter: Gao Shang (Tsinghua University)
Author(s): Shang Gao;Zhonghao Fu;Liangjun Su;Xia Wang

This paper studies the estimation and variable selection in conditional factor models with high-dimensional instruments, where the coefficient matrix exhibits a low-rank and row-sparse structure. We propose a multi-stage estimation procedure that combines nuclear norm regularization and adaptive group LASSO regression to consistently estimate latent factors and row-sparse loading coefficients, while selecting relevant instrumental characteristics. We establish theoretical results for estimation consistency, selection consistency, and post-LASSO inference for estimators of factors and loading coefficients at multiple stages. Furthermore, we implement a singular value thresholding procedure to determine the number of factors. Simulation results demonstrate the effectiveness of our estimators in consistently estimating factor loadings, selecting the appropriate number of factors, and conducting inference. Finally, we apply the proposed method to an empirical study on asset return prediction, showcasing its practical utility in real-world applications.

Forecasting of Integer-valued Time Series

Presenter: J Yuan (The University of Manchester)
Author(s): Jingsong Yuan;Xixi Li

We consider the forecasting of integer-valued time series using a Gaussian copula model with discrete marginals, which has a latent time series satisfying a vector autoregressive (VAR) model. We take a semi-parametric approach and allow the marginal distributions to be unspecified. The dependence parameters in the Gaussian copula/VAR model will be estimated by either Monte Carlo Expectation Maximisation (MCEM) or directly maximising an approximation of the log-likelihood function. Both methods involve importance sampling from truncated multivariate normal distributions using a variation of the GHK simulator (Geweke, Hajivassiliou and Keane). Forecasting from fitted models will be carried out using simulated samples and results will be compared with existing methods such as ETS, VAR and tscount.

Deep Learning for Time-Varying Probabilistic Forecast Combination of Intermittent Time Series

Presenter: Xixi Li (Tsinghua University)
Author(s):

Authors: Jian Zong, Xixi Li

Abstract: Forecasting intermittent time series presents significant challenges in many real-world applications. This paper proposes a novel approach for time-varying probabilistic forecast combination using deep learning techniques to address these challenges. Empirical experiments on real-world retail demand data show that our method outperforms traditional models, including time-invariant combinations and simple averaging techniques.

\subsection{The Genestrapping®-Method for less Inventories and generally 100% Customer Service}Presenter: Klaus Spicher (Retired from TH OWL, Germany)
Author(s): Klaus Spicher

Genestrapping® represents an innovative approach for Inventory Management without forecasting in the S&OP-Environment. The basic idea behind the method is an analogy of the biological inheritance process applied on time series. Genestrapping® succeeds in generating generally 100% Customer Service and near to minimum Inventories, while omitting traditional Forecasts, safety stock etc…

Shrinkage Estimators for Vector Exponential Smoothing

Presenter: Kandrika Pritularga (Lancaster University)
Author(s): Kandrika Pritularga;Ivan Svetunkov;Nikolaos Kourentzes

Vector Exponential Smoothing (VES) is a multivariate extension of Exponential Smoothing. Like many other multivariate models, it suffers from overparameterisation and overfitting especially when the sample size is limited. These consequently affect the forecasting performance of the model. We investigate shrinkage as a solution to these issues and highlight that forcing some parameters to zero arbitrarily is an ad-hoc form of shrinkage. We propose a shrinkage estimator for VES to shrink not only the smoothing parameters but also the covariance matrix, as both compensate for each other. We conducted simulation experiments to test the efficacy of the proposed estimator. The findings show that univariate Exponential Smoothing models outperform the various VES models in a multivariate setting, and the ad-hoc shrinkage performs better than any proposed shrinkage estimators. We suggest exploring the properties and the behaviour of the loss function of Exponential Smoothing and VES.

A Spatial-temporal Attention Based BiLSTM for Power Grid Investment Prediction

Presenter: Yijing Wang (China University of Petroleum (East China))
Author(s): Yijing Wang;Xiaofeng Xu

Mining the investment potential of the power grids in countries along the Belt and Road Initiative (BRI) is crucial for global energy transition. However, significant differences in national conditions make it hard to capture correlations among influence features. Moreover, policy and technology changes during the energy transition cause non-linear and dynamic changes in feature importance over time, further complicating the prediction of investment potential. To address the above-mentioned issues, this paper proposes a Bidirectional Long Short-Term Memory network incorporating the double-layer Attention Mechanism (AM-BiLSTM). In this model, the spatial and temporal attention mechanisms are respectively used to adaptively mine the feature correlations among different countries and moments, and assign differential weights to different features and time steps. The BiLSTM neural network is utilized to fit the data and construct the prediction model. The experiments showed that the proposed model outperforms other benchmark and widely-used models in the task of predicting the investment potential. The results can provide a reference for investment decisions in the energy and power industry in countries along the BRI.

Enhancing Interpretability and Accuracy in Exchange Rate Prediction with Explainable Artificial Intelligence and Causal Discovery

Presenter: Siran Fang (School of Economics and Management, University of Chinese Academy of Sciences, Beijing 100190, China)
Author(s): Siran Fang;Yunjie Wei;Shouyang Wang

Deep neural networks are widely used for exchange rate forecasting, but their “black-box” nature limits interpretability. This study integrates deep neural networks with Explainable Artificial Intelligence (XAI) and the Peter and Clark Momentary Conditional Independence (PCMCI) method to enhance both prediction accuracy and model transparency. We employed a Long Short-Term Memory (LSTM) network to forecast the exchange rates of the Euro, British Pound, Japanese Yen, and Chinese Yuan against the US Dollar from January 2004 to March 2024. SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) provided global and local insights into the LSTM model, while PCMCI identified causal relationships between exchange rates and other variables. Results show that exchange rate itself is crucial for short-term predictions, supporting the Efficient Market Hypothesis. Unique economic factors, including the US Dollar Index, copper prices, Nikkei 225 Index, and bank reserves, are key predictors for each currency. Financial market volatility, policy uncertainty, geopolitical risks, and

risk aversion sentiments drive exchange rate fluctuations during major shocks like financial crises and geopolitical conflicts. XAI-based feature selection validated the model's interpretability and improved predictive performance. This research offers a novel approach to understand exchange rate dynamics, enhancing forecast accuracy and providing insights for policymakers.

The Instability of Leading Indicators in Forecasting Austrian Inflation: Lessons from the COVID-19 Pandemic and the Energy Crisis

Presenter: Friedrich Fritzer (Central Bank of Austria)
Author(s): Friedrich Fritzer

This analysis tests 25 macroeconomic indicators for their ability to predict Austrian HICP inflation and evaluates three methods of combining these indicators into a composite forecast. The key findings are: First, for the evaluation period from early 2007 to the fourth quarter of 2023, competitors' import prices, oil prices and domestic output prices for consumer goods are found to be the best individual leading indicators across various forecasting horizons (one, four and eight quarters ahead). Second, indicator performance varies over time. The forecasting performance of the output gap, for instance, declined considerably during the COVID-19 pandemic and the energy crisis, while that of other indicators like oil prices and competitors' import prices on the domestic market improved. Third, for the period before 2020, composite indicators produced better forecasts than individual indicators over the entire forecasting horizon. This no longer holds when we include the pandemic and the energy crisis in the evaluation period. Then, two of the top three individual indicators, namely competitors' import prices on the domestic market and domestic output prices for consumer goods, outperform combined indicators over the medium- and longer-term horizon (four and eight quarters ahead). Fourth, both individual and composite indicators outperformed autoregressive forecasts, especially in medium- and long-term predictions.

Multimodal Aspect-Based Sentiment Analysis Meets Deep Learning for Tourism Forecasting

Presenter: Xin Li (University of Science and Technology Beijing)
Author(s): Xin Li;Xin Li;Shouyang Wang;Xu Zhang

The increasing availability of textual user-generated content has significantly advanced tourism demand forecasting; however, the extraction and interpretation of sentiment from multimodal data remains a major analytical challenge. Recent advancements in large language models enable more systematic annotation and structuring of aspect-level sentiment features from vast volumes of social media content. To address these challenges, this study proposes a multimodal aspect-based framework that integrates natural language processing, computer vision, and deep time-series forecasting to enhance the predictive accuracy of tourism demand models. Specifically, (1) we develop domain-adapted BERT and Vision Transformer models to classify sentiment in 20,509 Weibo posts and 110,871 associated images from a Chinese tourist destination, thereby capturing the complementary roles of textual and visual social media data; (2) we employ a ChatGPT-4o-driven large language model for annotation and optimization to extract fine-grained aspect-opinion-sentiment triples from textual data, while leveraging DeepSentibank to identify adjective-noun pairs representing visual affective semantics; and (3) we integrate the resulting aspect-level sentiment features as exogenous variables into three state-of-the-art deep learning forecasting architectures—Informer, FEDformer, and PatchTST—to evaluate their contribution to tourism demand prediction. Empirical results demonstrate that incorporating aspect-based sentiment features significantly improves forecasting accuracy compared to models relying solely on historical visitation data. Notably, we identify a "sentiment-consistency effect," where aligned sentiment polarity across text and images is associated with lower variance in tourism forecasts. By integrating multimodal sentiment analysis with deep learning-based forecasting, this study contributes to the methodological advancement of tourism demand prediction and provides actionable insights for destination managers through real-time sentiment monitoring and scenario-based decision support. The proposed framework offers a transferable approach to data-driven tourism management, bridging the gap between computational social science and practical forecasting applications.

Complex Data Modeling and Analysis with Applications

Presenter: Shaolong Sun (Xi'an Jiaotong University)
Author(s): Shaolong Sun;Haoqiang Sun

Tourism vlogs leverage multimodal integration of narrative and visual elements to create immersive destination experiences, serving as vital tourism marketing tools. However, mechanisms between multimodal elements and video appeal remain underexplored. Anchored in Expectation-Disconfirmation Theory, this study examines the effects of video scene similarity, audio scene similarity, and audiovisual consistency on vlog attractiveness. Through deep learning and econometric analysis of 466 global TikTok tourism vlogs, results indicate video and audio similarities negatively affect attractiveness, while audiovisual consistency positively enhances engagement. The synchronization of emotional expression and rhythm between visuals and narrative significantly improves viewing experiences. This research advances theoretical understanding of vlog marketing and provides actionable insights for destination managers in optimizing video strategies.

Daily Tourist Flow Forecasting Considering Short-Video Visual Features

Presenter: Chengyuan Zhang (Xidian University)
Author(s): Chengyuan Zhang

The rise of short-video platforms has reshaped the paradigm of tourism information dissemination, creating a closed-loop behavioral chain from "online visual engagement to offline check-in feedback" that drives tourist decision-making. However, the existing approaches to predicting tourism demand that incorporate predictive short-video visual variables still leave room for improvement. To address this, this study constructs a Scenic Short-video Search Index (comprising keyword and content search volumes) to indicate tourist interest trends, and a Composite Popularity Index (encompassing video publication counts, playback volumes, and dissemination metrics) to capture the overall heat of short videos. Based on models such as MLP, LSTM, and GRU, we propose a tourist flow forecasting framework that integrates these visual variables. The results show that incorporating both the Short-video Search Index and the Composite Popularity Index significantly improves model accuracy. Specifically, when these indices are introduced into the MLP, LSTM, and GRU models, the MAE decreases by 3.4%, 8.8%, and 19.1% respectively, and the MAPE drops by 5.6%, 4.3%, and 14.0%. These findings demonstrate the effectiveness of short-video visual variables in predicting tourist flows, providing new data dimensions and optimization pathways for tourist flow forecasting.

Forecasting Financial Reporting Fraud in Financial Companies

Presenter: Nurul Hasanah Uswati Dewi (Universitas Hayam Wuruk Perbanas)
Author(s): Nurul Hasanah Uswati Dewi

This research aims to forecast the significance of the Fraud Hexagon Theory in explaining fraudulent financial reporting in companies. The theory encompasses the variables stimulus, capability, collusion, opportunity, rationalization, and ego. The sample consisted of financial sector companies listed on the Indonesia Stock Exchange (IDX), selected using purposive sampling techniques. The findings, based on data processing and the theoretical framework, indicate that Stimulus (proxied by financial targets), Capability (proxied by change in director), Collusion (proxied by cooperation with government projects), Rationalization (proxied by change in auditor), and Ego (proxied by the frequent number of CEOs' pictures) do not significantly affect fraudulent financial reporting. However, Opportunity (proxied by the nature of the industry) does have a significant effect. Logistic regression analysis was employed in this research.

Monitoring Consumption Vulnerability with Economic News and Macroeconomic Big Data

Presenter: Shuting Liu (Shanghai University of Finance and Economics)
Author(s): Shuting Liu;Shiqi Ye;Guangfei Sun

We developed a novel consumption-at-risk (CaR) framework that integrates economic news and macroeconomic big

data to monitor consumption vulnerability in China. We first extract interpretable economic topics from over 1.7 million articles published in mainstream Chinese newspapers using the online latent dirichlet allocation (oLDA) topic model. Then, a mixed frequency quantile regression model is combined with various regularization techniques, in order to examine the contribution of large-scale news topics and macroeconomic indicators to monitoring consumption risks under the CaR framework. Empirical results demonstrate that news data provide incremental information value in monitoring CaR, particularly at the lower quantiles. In comparison with traditional macroeconomic indicators, news data characterized by broader information coverage and higher update frequency, significantly enhance the accuracy and timeliness of CaR monitoring. A further analysis of indicator relative importance suggests that macroeconomic conditions, financial market dynamics, labor market status, and international relations play critical roles in shaping consumption vulnerability.

Business Cycle Dynamics after the Great Recession: An Extended Markov-Switching Dynamic Factor Model

Presenter: Laurent Ferrara (SKEMA Business School)
Author(s): Laurent Ferrara;Catherine Doz;Pierra-Alain Pionnier

As illustrated by the Great Recession and the Covid pandemic, macroeconomists need to account for sudden and deep recessions, as well as shifts in macroeconomic volatility and trends in GDP growth. In this respect, we put forward an extended Markov-Switching Dynamic Factor Model by incorporating switches in volatility and time-variation in trend GDP growth. We empirically show that volatility switches improve the detection of business cycle turning points, that the Great Recession led to only a temporary increase in volatility, and that the US trend GDP growth has been declining since the early 2000s. Information criteria, marginal likelihood comparisons and improved real-time performance support our proposed model. Our results show that the model accommodates well the pre- and post-Covid times. Real-time density forecasts show that our model is able to capture the growth-at-risk main features, namely the left-skewed distributions during recessions.

A New Real-time Composite Indicator for German SME Business Cycles Based on Taxes and Balances

Presenter: Dirk Beinert (DATEV eG)
Author(s): Dirk Beinert

Current business cycle reports on the German economy are often constrained by a very limited sample size or a significant publication delay of several months. Nevertheless, scientific and administrative institutions claim to forecast economic developments months or even years into the future. This presentation examines the leading economic indicators and reports in Germany, assessing their predictive capabilities. Additionally, it introduces the development of a new composite indicator based on real-time tax and financial data from a substantial share of German SMEs, as transmitted via tax consultant software.

Forecasting the Demand for Recycled Plastics using a Digital Twin for Regional Plastic Cycles

Presenter: Christian Menden (Technical University of Applied Sciences Würzburg)
Author(s): Christian Menden;Ulrich Müller-Steinfahrt

Plastic waste that accumulates in manufacturing companies, e.g. as production residues, packaging material or transport carriers, is very often not a material that can be utilized economically or in terms of resource efficiency. This means that large quantities of plastic waste are often disposed of in waste incineration plants for energy recovery. In Germany, due to the lack of a standardized recycling system for industrial plastic waste - like the dual system for consumer waste - the valuable materials do not enter the cycle and are permanently lost. On the other hand, companies in the recycling industry that process post-industrial waste are desperately looking for high-quality plastics that are reliably and predictably available as input material. Against the background of increasingly strict regulatory requirements (e.g. recycled material usage quotas, plastic tax, CO2 footprint) and the demands of industry and consumers, processors are increasingly looking for new approaches to establish secondary plastics into their products.

Yet, there exist many practical obstacles to establishing a circular network flow of industrial plastics, e.g. lack of trust regarding the quality and a reliable availability of recycled materials, lack of data on the processing history and other quality parameters, complex production processes and interdependencies with high uncertainty at multiple steps in the envisioned circular network, usually small batches of recycled materials that make logistic concepts for pick-up, collection and distribution of recycled materials not profitable, etc. To realize a circular system for industrial plastics, we launched a pilot project "Cooperative Plastics Cycles Mainfranken" (a regiopolis region in Bavaria, Germany) with multiple research and industry partners. We identified the accumulated plastics streams on a weekly basis within the network and analyzed possible symbiotic relations. As next, we develop a digital twin of individual production processes, determine the most relevant data and set up an infrastructure for the efficient transfer and storage of digital sensor data. We model the demand and availability of plastic waste for different processes and analyze the similarity of uncertainty estimates to determine the potential for symbiotic flows and the effectiveness of different logistic concepts.

Forecasting the Forecast: Bridging Prediction and Reality in Automotive Supply Chains

Presenter: Andreas Rükgauer (THWS Business School)
Author(s): Andreas Rükgauer;Meghashree Muralidhar

Reliable demand forecasting is necessary for optimizing operations along the automotive supply chain. A peculiarity of automotive supply chains is the fact that alongside with the call-offs for shipments, a long term forecast is provided. In this study, real data from an automotive production part supplier is used daily shipments and daily forecasts 180 days out. Past research on this data and also in related work has exploited the past actual shipment information only, treating it as a time series and providing a future prognosis from corresponding models. This study tries to enhance the prediction accuracy of the shipment by utilizing neural network models to cover the gap between forecasted shipment data and actual shipment volumes. A key breakthrough in the current approach is the ability to forecast actual shipments 15 days into the future rather than extending the existing forecasts. The data was extensively preprocessed, which involved structuring, normalizing and eliminating inconsistencies such as weekend and vacation gaps to ensure model robustness. To optimize predictions, a comprehensive experiment cycle is designed to analyze hyperparameter configurations, model architecture and performance The study hints that the current approach helps in reducing forecasting errors and improves the final predictions. The work could be used to enhance logistics and demand planning in supply chain environments, further laying a solid groundwork for better future developments with neural network driven solutions for better predictive modelling system.

The Structure of China's Ocean Economy and Its Future Trends Projection

Presenter: Zengkai Zhang (Xiamen University)
Author(s): Zengkai Zhang

The ocean economic input-output (IO) tables are essential databases for analyzing the ocean economy. However, the IO tables issued by Chinese government's statistical department adopt the conventional economic sector classifications, which fall short in capturing the heterogeneity of the ocean and land sectors. To fill this gap, we used the stripping coefficient method to construct the ocean economic IO tables of 11 coastal provinces in China, relying on the latest official provincial IO tables in 2017 and ocean economy data published by official institutions. Based on provincial ocean economic input-output tables, we analyze the structural composition of China's ocean economy. Finally, we project its future development trends.

Estimating and Predicting the Effect of China-US Trade Friction on Bunker Fuel Market volatility

Presenter: Xiaoxia Li (Shanghai Maritime University)
Author(s): Xiaoxia Li;Tse Leung Yip

This paper investigates how these geopolitical tensions have affected the volatility of the bunker fuel market, a critical component of maritime transportation. The China-US trade friction, initiated in 2018, has reshaped global

trade patterns, disrupted supply chains, and influenced commodity markets and maritime transportation. We adopt the synthetic control method to estimate and predict the effect of China-US trade friction on bunker market volatility in several ports of China and US. Limited by the access to data, it is difficult to find the best control groups for bunker markets. This study solves the difficult issue by constructing the synthetic markets as control units. And then we compare bunker markets with the synthetic bunker markets to obtain the treatment effect of bunker markets. The method adopted by this study uses the data-driven procedure to construct comparison groups, which reduces the non-transparency in selecting comparison control groups. The findings suggest that the China-US trade friction makes more fluctuated volatility in China and US ports, which would harm the interests of relative stakeholders and investors. Also, our results forecast the trends of bunker markets volatility, and then identify the significant linkages between trade policy shifts and market instability while projecting bunker market volatility under trade friction scenarios. The findings will provide actionable insights for policymakers, shipping firms, and energy traders to mitigate risks and optimize decision-making in an uncertain geopolitical climate.

A Novel Shipping Freight Rate Forecasting Framework based on LLM

Presenter: Jinyu Wu (Shanghai Maritime University)
Author(s): Jinyu Wu;Yanhui Chen

In order to explore the application mode and prediction ability of large language model in the field of shipping freight rate prediction, this paper constructs an analysis framework based on natural language processing technology, and conducts text mining on the professional reports of the dry bulk freight rate market from 2020 to 2025, and applies the semantics obtained from the analysis to the prediction of the Baltic dry bulk index (BDI). In this study, the original news is preprocessed by deepseek-r1 model, and then a two-stage classification strategy is adopted: the initial classification is based on the title similarity, and then the text content is vectorized by gte-base-en-v1.5 model. The high-dimensional text embedding is reduced to nine main dimensions through principal component analysis (PCA), which can effectively capture the semantic characteristics of the report content. Finally, the mapping relationship between text semantic features and BDI is established. The experimental results show that there is a significant correlation between the specific semantic pattern in the report and the subsequent market price changes, which provides an application way for the prediction of shipping freight rates using the large language model, and also provides a new perspective for understanding the impact of information dissemination on the dry bulk market. The purpose of this study is to explore the correlation between the Research Report of dry bulk trade industry and the change of market index. This study provides a quantitative method based on text analysis for investment decision-making and market prediction, and has practical application value for improving the accuracy of financial market prediction.

Global Banks′ Macroeconomic Expectations and Credit Supply

Presenter: Xiang Li (Hall Institute for Economic Research)
Author(s): Xiang Li

We investigate how global banks' macroeconomic expectations for borrower countries influence their credit supply. Utilizing granular data on varying expectations among banks lending to the same firm at the same time, combined with an instrumental variable approach, we find that more optimistic GDP growth expectations for a borrower country are strongly linked to increased credit supply. Specifically, a one standard deviation increase in a lender's GDP growth expectation for a borrower country corresponds to an increase of approximately $75.35 million in the bank's lending to that country. In contrast, global banks' short-term inflation expectations do not show a significant impact on their credit supply.

Forecasting Japanese Recessions Using Machine Learning and Mixed-Frequency Data

Presenter: Yusuke Oh (Bank of Japan)
Author(s): Yusuke Oh;Mototsugu Shintani

This study investigates the forecasting of Japanese recessions. We propose a framework based on machine

learning models that accommodates mixed-frequency data. In addition to standard macroeconomic variables, we consider a broad range of predictors that the literature has found useful for recession forecasting, including government bond term spreads, financial variables such as the Debt Service Ratio (DSR), and text metrics extracted from newspapers—the latter of which have been less studied in this context. A pseudo-real-time, out-of-sample forecasting exercise was conducted over the past 30 years, examining comprehensive combinations of predictors to assess the additional benefits provided by each category. The results show that machine learning models substantially outperform traditional logit benchmarks across short-, medium-, and long-term horizons. Model selection via the Model Confidence Set (MCS) procedure demonstrates that text-based indicators and financial variables effectively enhance short-term forecasting performance, whereas term spreads emerge as effective predictors for long horizons. Our findings also indicate a limited advantage of using higher-frequency data, which may be attributed to the presence of abundant variables that already capture recession-related information. We further apply SHapley Additive exPlanations (SHAP) to our machine learning models and extract interpretable sparse principal components from the text metrics. This analysis reveals that different recession episodes in Japan can be characterized by distinct narrative profiles. For instance, some recessions show a dominant influence of corporate distress signals—such as bankruptcies and non-performing loans—while others are more closely linked to financial turmoil or deflationary pressures. These insights not only enhance our models' interpretability but also offer a richer understanding of the heterogeneous economic drivers behind each recession.

A Flexible and Collaborative Method for Generating Extreme Macro Scenarios

Presenter: Ruoshi Shi (Industrial and Commercial Bank of China postdoctoral research workstation)
Author(s): Ruoshi Shi

With the global economy entering a downward phase, China's economic and financial markets are facing unprecedented challenges. The rapid evolution of financial risks makes it difficult for traditional means to quickly assess and accurately predict risks, so it is urgent to study and develop forward-looking risk assessment methods. This paper focuses on how to respond flexibly and collaboratively to sudden downside risks in the financial system under the challenges of new risk situations, and proposes a more flexible and collaborative stress scenario modeling and generation method. As the complexity of risks increases and proactive regulation control enhances, the existing risk models suffer from problems such as increasing unmeasured variability, difficulty in dimension expansion and insufficient synergistic pressure. In this paper, first of all, constructed a flexible and collaborative exogenous driven shock to make the exogenous shock embedded in the state-space model, which overcomes the difficulty of model compression and improves the consistency of impact amplitude among model groups. Secondly, through comprehensive analysis of risk sources, expansion of core scenario indicators and orderly nesting and segmentation, we built a linkage pressure scenario set to facilitate the transmission of pressure shock among multiple scenario sets. In addition, a pulse drive method based on prediction variance decomposition is proposed to flexibly apply pressure to endogenous and exogenous pressure shocks and improve the synergy of driving pressure in the model. This work has shown good practicability and applicability in the actual stress testing work of commercial banks and regulatory institutions.

Nowcasting US Economic Activity in a Changing World and a Data-Rich Environment

Presenter: Chen Ziyang (Xiamen University)
Author(s): Ziyang Chen;Shiqi Ye;Tingguo Zheng

The rapid evolution of the global environment has intensified volatility and recurrent structural shocks, making the nowcasting of economic activity increasingly challenging. At the same time, the growing availability of big data offers new opportunities to address these challenges. This paper proposes a fast estimation method for a time-varying parameter mixed-frequency dynamic factor model (TVP-MFDFM), which accounts for potential structural changes in the economy while leveraging big data to enhance nowcast accuracy. Specifically, we utilize the bilinear structure within the time-varying factor model, combining Taylor expansion and a classical mixed-frequency specification to develop a simple and efficient estimation method for the TVP-MFDFM. This allows high-dimensional and high-frequency data to be effectively utilized for economic activity nowcasting. Empirical results demonstrate that: (1) the introduction of time-varying factor loadings improves the robustness of model estimation, leading to stable

economic activity measures; (2) the mixed-frequency setup enhances out-of-sample forecasting performance by efficiently incorporating high-frequency data; (3) the integration of big data significantly improves the effectiveness of economic activity nowcasting. Furthermore, the proposed model shows substantial advantages over benchmark models and professional forecasters in numerical exercises and real data evaluations.

Nowcasting Macroeconomic Downward Risk: The Role of Time-Variation, Skewness, and Group-Sparsity

Presenter: Mengdi Lv (University of Science and Technology of China)
Author(s): Mengdi Lv;Shiqi Ye

Rapid changes in the global environment have significantly increased economic uncertainty, leading to growing attention on nowcasting macroeconomic risks, particularly the tail risks of economic growth, named as growth-at-risk. This paper integrates mixed-frequency, group-sparsity, and a score-driven skewed t-model specification into a unified framework, utilizing real-time high-dimensional data to update the time-varying location, scale, and skewness parameters in the model, thus capturing the time-variation of tail risk of economic growth. A Bayesian estimation method is developed to enable real-time nowcasting of tail risks in US GDP growth, while also enabling dynamically extracting important predictive variables. Empirical results demonstrate that: (i) the introduction of time variation, skewness, and mixed-frequency significantly improves the out-of-sample nowcasting performance of US growth-at-risk; (ii) high-dimensional information together with group-sparsity significantly improves the real-time nowcasting of location, scale, and skewness parameters, with Fred-MD showing strong predictive power for scale and skewness, but weaker performance for location; (iii) effective predictive variables vary across different economic environments, and the proposed model can select variables from big data that effectively predict US growth-at-risk, thereby enhancing forecasting accuracy; (iv) variables predictive of US growth-at-risk exhibit clear time-varying characteristics.

When Large Language Model Meets Textual Data: Toward Narrative-Driven Macroeconomic Forecasting

Presenter: Yue Wu (Academy of Mathematics and Systems Science)
Author(s): Yue Wu

This paper introduces a novel approach that integrates large language model for the timely extraction of economic signals from textual data, referred to as the targeted-LLM (TLLM) method. The proposed methodology combines the strengths of two popular approaches: topic modeling and sentiment lexicon, in textual data extraction and macroeconomic forecasting. By introducing a well-designed zero-shot prompt templates, the TLLM method enables the extraction of multidimensional information from economic news, including news topics, impact strength, impact direction, and keywords, specific to a target predictive variable. Furthermore, We apply the TLLM method to forecast and nowcast the US GDP growth rate, inflation rate, and unemployment rate. The in-sample results demonstrate that the TLLM method accurately and comprehensively extracts useful information related to specific macro indicator from economic news, with superior interpretability. The out-of-sample forecasting results show that the topic-related indicators obtained from the TLLM method significantly outperforms other competitors, especially the online latent dirichlet allocation (oLDA) approach for all the indicators considered, and provides added value compared to FRED-MD in the extracted information from economic news. Out-of-sample nowcasting results reveal that the two impact strength indicators constructed from the TLLM method offer more timely and accurate nowcasting performance than existing sentiment indicator. The proposed method is highly replicable, operationally practical, and capable of self-upgrading as LLMs evolve, providing valuable insights for future research in the relative domain.

The Power of Themes: A Topic-Aware Approach to Predicting Engagement with Tourism Destination Videos

Presenter: Peijun Shi (University of Science and Technology Beijing)
Author(s): Peijun Shi;Xin Li;Xiangbin Yan

In the era of short-form video dominance, understanding the drivers of user engagement with tourism destination content has become increasingly critical for effective digital communication and destination branding. While prior research has acknowledged the role of content features in shaping user behavior, the specific influence of thematic features—particularly in cultural heritage contexts—remains underexplored. This study addresses this gap by proposing a topic-aware framework to predict user engagement with destination videos, using the official Douyin account of the Palace Museum as a case study. Integrating multimodal data and machine learning techniques, we systematically examine how themes extracted from video texts affect user interactions, including likes, comments, favorites, and shares. To extract high-level thematic signals, we apply pre-trained large language models for topic classification and sentiment analysis of video captions. These features are then combined with visual and audio attributes to construct predictive models across multiple engagement types. Results show that cultural and historical themes—especially those emphasizing national heritage, emotional storytelling, and educational value—exhibit strong predictive power in driving user engagement. Further, SHAP-based interpretability analysis reveals nonlinear and heterogeneous effects of thematic content, highlighting the mediating role of perceived cultural identity and emotional resonance. Compared with traditional content features, thematic signals offer higher semantic granularity and contextual depth, enabling more accurate behavioral prediction. The study also demonstrates that not all themes equally influence all types of user behavior; for instance, emotionally charged historical narratives tend to drive sharing, while aesthetically rich cultural displays are more likely to attract likes and comments. These findings offer both theoretical and practical contributions. Theoretically, the study expands the understanding of multimodal engagement prediction by incorporating thematic dimensions rooted in cultural communication. Practically, it provides actionable insights for destination marketers to design more resonant and targeted video content. By emphasizing the power of themes, this research underscores the strategic value of culturally aligned storytelling in enhancing digital destination communication effectiveness.

Can AI Help Tourists Predict the Reasonable Prices of the Hotels? Decomposing the Gap between Hotels′ price and Tourist Perceived Value.

Presenter: Zhao Lu (China University of GeoSceicnes(Beijing))
Author(s): Lu Zhao;He Zhu;Shouyang Wang;Mingli Zhang

Nowadays, tourists resort to social media to search for the hotel information. According to the value - price theory, customers who develop the high perceived value tends to reserve the hotels, otherwise they reject the hotel. Hence, it is very important for hotels to comprehend the customers'perceived value in online context. However, due to the information uncertainty and brand feature, tourists are faced with the challenge of gaining the reasonable price of the hotels, leading to the price premium and tourists complain. To solve this issue, we use multi-modal methods to visibly measure the gap between the perceived value and price. We first crawel the tourist generated images and texts on Rednote and ctrip.com platform. Then we use ViLBERT model to extract the key visual and the textual information from the reviews. By collecting the semantic feature, we construct the tourists perceived value dictionary: social value, function value, emotional value and symbol value. Besides, we utilize the two-way fixed effect model to compute brand effect. We also use the XGBOOST methods to connect the perceived value and hotels prices. Our findings can help hotel managers intuitively notice the gap between the list price and actual tourists perceived value, which could facilitate marketers to better comprehend the tourists from online content.

ANY LEARNING OBJECTIVE: ADAPTING THE FOUNDATION FORECASTING MODELS TOWARDS YOUR BUSINESS GOALS AND CONSTRAINTS

Presenter: Geoffrey Negiar (The Forecasting Company)
Author(s): Geoffrey Negiar

Generating accurate probabilistic forecasts is a critical yet challenging task in many applications.Traditional probability estimation techniques, such as likelihood maximization and maximum a posteriori estimation, often lead to a misalignment between the model's learning objective and the actualevaluation metric used in practice. This misalignment becomes even more problematic in complexforecasting scenarios that require specific constraints, such as positivity or hierarchical coherenceacross levels.In this talk, we introduce Coherent Learning Objective Reparameterization (CLOVER) (Olivareset al., 2024), a technique that enables probabilistic models to produce samples that are

differentiablewith respect to the model parameters (Kingma & Welling, 2014). This allows optimization onarbitrary differentiable learning objectives and constraints, ensuring alignment with the goals ofthe forecasting system. We demonstrate CLOVER's effectiveness on hierarchical forecasting tasks,achieving state-of-the-art performance in six public benchmark datasets, with improvements in bothprobabilistic and mean forecast accuracy.

ANY HORIZONS: LOOKING INTO THE DISTANT FUTURE

Presenter: Malcolm Wolff (Amazon)
Author(s): Malcolm Wolff

Varying horizons still pose a challenge for current architectures. Neural forecasting models typicallyrely on either recurrent or joint decoders for generating multi-horizon-ahead forecasts. However,both strategies come with inherent limitations. Recurrent decoders are computationally expensivedue to their sequential nature, which prevents parallelization and makes them prone to error accumulation over long horizons (Marcellino et al., 2006). Meanwhile, joint decoders tie the modelparameter count to a fixed forecast horizon, limiting flexibility when different horizon lengths arerequired (Amir & Souhaib, 2016).In this talk, we introduce SPADE, a highly efficient decoder-only architecture that replaces horizonspecific parameters with a horizon-positional encoding, providing an alternative to traditional recurrent and joint forecast strategies. Our new forecast strategy enables parallelization, reduces theaccumulation of long-term errors, and overcomes the rigidity of the joint decoders. Furthermore,we demonstrate that SPADE's decoder improves forecast robustness, particularly against carry-overeffects in time series with large peaks (Wolff et al., 2024). Finally, we show that previously observedaccuracy trade-offs in horizon-independent forecasting (Bao et al., 2014) are significantly mitigatedwhen trained on sufficiently large datasets.The SPADE model currently serves in production forAmazon demand forecasting.

SPECIALIZED SPARSE FORECASTING MODELS

Presenter: Hanjing Zhu (Amazon)
Author(s): Hanjing Zhu

Despite significant advancements in time series forecasting, accurate modeling of time series withstrong heterogeneity in magnitude and/or sparsity patterns remains challenging for state of the artdeep learning architectures. We identify several factors that lead existing models to systematicallyunder-perform on low magnitude and sparse time series, including loss functions with implicit biases toward high-magnitude series, training-time sampling methods, and limitations of time seriesencoding methods. To address these limitations, we introduce SPADE-S Wolff et al. (2025), a robustforecasting architecture with a novel multi-head convolutional encoder and a model arm specifically designed to handle sparse multi-variate time series. SPADE-S significantly reduces magnitudeand sparsity-based systematic biases and improves overall prediction accuracy; empirical resultsdemonstrate that SPADE-S outperforms existing state-of-the-art approaches across a diverse set ofuse-cases in Amazon's online retail demand forecasting. In particular, we show that on the mostaffected time series, SPADE-S can improve forecast accuracy by up to 15%. This results in P90overall forecast accuracy gains of 2.21%, 6.58%, and 4.28%, and P50 forecast accuracy gains of0.92%, 0.77%, and 1.95% respectfully, for each of three distinct datasets, ranging from 3 million to700 million series, from a large online retailer.

Forecasting the Direction of Changes in Bitcoin Quotes based on Turning Point Patterns Identified by Dynamic Time Warping

Presenter: Michal Bernardelli (SGH Warsaw School of Economics)
Author(s): Michal Bernardelli;Bartosz Witkowski;Mariusz Kozakiewicz

Theoretical background: Automating investor actions on cryptocurrency exchanges is necessary due to their specificity. Research in this area includes arbitrage opportunities, predicting the direction of changes in cryptocurrency quotes, and constructing investment strategies. However, one of the critical aspects of analysing investment instrument quotes is identifying turning points and trying to find regularities in the behaviour of the time series of these quotes. Purpose of the article: The study aimed first to determine the degree of similarity of the periods

preceding the turning points of Bitcoin's quotations. The choice of this cryptocurrency was related to its dominance in the cryptocurrency market and its proven impact on altcoin prices. The second set goal was to verify the possibility of developing a pattern reflecting peaks and troughs in the cryptocurrency market near the reversal of the Bitcoin quotes trend. The third goal is to use the developed patterns to forecast the upcoming change in the direction of prices. Research methods: This paper presents an approach to analysing the similarity between different phases of business cycles based on the concept of dynamic time warping (DTW). The DTW measure was the basis for cluster analysis and was used to determine the distance between clusters composed of time series. An integral part of the study was proposing a procedure for constructing patterns of quotations' behaviour in a specified window before the identified turning point. Those patterns are the basis of the forecasting procedure. Main findings: An empirical analysis was conducted using the example of Bitcoin. A more significant similarity was observed between the contraction phases in comparison to the expansion phases. In the case of contraction phases, two types of patterns dominate, while in the case of expansion phases, there is one dominant pattern. Based on the analysis of the dominant patterns, it can also be stated that before the reversal of the situation in the cryptocurrency market, stable and systematic increases and decreases occurred more often than quotation patterns with high volatility. Pattern analysis allows for the construction of forecasts, which should be treated as an indicator of a trend change.

Hybrid Models for Financial Forecasting: Combining Econometric, Machine Learning and Deep Learning Models

Presenter: Dominik Stempień (University of Warsaw)
Author(s): Dominik Stempień;Robert Ślepaczuk

Financial time series forecasting plays a fundamental role in investment decision-making, risk assessment, and the development of trading strategies. Econometric time-series models, often struggle with the complexity of financial markets due to their low signal-to-noise ratio and inherent uncertainty. On the other hand, machine learning (ML) and deep learning (DL) methods have proven effective in identifying nonlinear structures in data. A less commonly used approach to time-series modeling consists of merging traditional methods and ML/DL techniques. The hybrid methodology is based on the assumption that combining diverse individual models allows to obtain more accurate forecasts. This study systematically develops and evaluates various hybrid modeling approaches by combining traditional econometric models (ARIMA and ARFIMA models) with ML and DL techniques (LSTM, XGBoost and SVM models) to forecast financial time series. The empirical analysis is based on two distinct financial assets: the S&P 500 index and Bitcoin. By incorporating over two decades of daily data for the S&P 500 and almost ten years of Bitcoin data, the study provides a comprehensive evaluation of forecasting methodologies across different market conditions. Furthermore, a novel time series validation method is employed to ensure robust model evaluation. A key aspect of this research is the comparative analysis of different hybrid model construction techniques, assessing their ability to capture both linear and nonlinear dependencies within financial time series. The results indicate that while hybrid models generally outperform individual statistical and ML/DL models, the choice of statistical and ML/DL components significantly affects forecast accuracy. Moreover, selecting the right hybridization method is crucial for forecasting efficacy. Notably, certain hybrid models enable the formulation of profitable trading strategies, surpassing the performance of the buy-and-hold approach. Additionally, a trading strategy based on a portfolio composed of both assets was tested to evaluate the practical applicability of the forecasts. By comparing different hybridization techniques, this research contributes to the advancement of financial forecasting methodologies. The study provides valuable insights into the effectiveness of hybrid models, offering a robust framework for both academic researchers and market practitioners seeking data-driven approaches to financial decision-making.

Informer in Algorithmic Investment Strategies on High Frequency Bitcoin Data

Presenter: Robert Ślepaczuk (University of Warsaw, Faculty of Economic Sciences, Department of Quantitative Finance, Quantitative Finance Research Group)
Author(s): Robert Ślepaczuk;Filip Stefaniuk

The article investigates the usage of Informer architecture for building automated trading strategies for high frequency Bitcoin data. Three strategies using Informer model with different loss functions: Root Mean Squared Error (RMSE), Generalized Mean Absolute Directional Loss (GMADL) and Quantile loss, are proposed and evaluated

against the Buy and Hold benchmark and two benchmark strategies based on technical indicators. The evaluation is conducted using data of various frequencies: 5 minute, 15 minute, and 30 minute intervals, over the 6 different periods. Although the Informer-based model with Quantile loss did not outperform the benchmark, two other models achieved better results. The performance of the model using RMSE loss worsens when used with higher frequency data while the model that uses novel GMADL loss function is benefiting from higher frequency data and when trained on 5 minute interval it beat all the other strategies on most of the testing periods. The primary contribution of this study is the application and assessment of the RMSE, GMADL and Quantile loss functions with the Informer model to forecast future returns, subsequently using these forecasts to develop automated trading strategies. The research provides evidence that employing an Informer model trained with the GMADL loss function can result in superior trading outcomes compared to the buy-and-hold approach.

Momentum, Reversal and Fundamentals: A Short-term Exchange Rate Forecasting

Presenter: Jiawen Ren (Zhongnan University of Economics and Law)
Author(s): Jiawen Ren;Yu Ren

This study introduces a functional-coefficient framework that integrates momentum effects, reversal patterns, and macroeconomic fundamentals for exchange rate prediction. To mitigate overfitting concerns in in-sample estimation, we employ an elastic-net regularization approach for model estimation, subsequently conducting rigorous out-of-sample forecasting analysis. Our empirical investigation yields three key findings: First, the proposed methodology demonstrates statistically significant outperformance over both random walk benchmarks and alternative prediction models in one-month-ahead forecasting across six major currency pairs. Second, decomposition analysis reveals that momentum indicators, reversal characteristics, and fundamental economic variables collectively contribute substantial predictive power. Third, implementing trading strategies based on our forecasts generates economically significant profits in government bond markets, confirming the model's practical utility in financial decision-making.

Integrating Wavelet Decomposition, Fundamental Economic Models, and Machine Learning Methods for Exchange Rate Forecasting

Presenter: Ruoxi Sun (Department of Economics, King's Business School, King's College London)
Author(s): Ruoxi Sun

Abstract This study proposes an exchange rate forecasting framework that integrates wavelet decomposition with machine learning and economic fundamentals models. Using redundant Haar wavelet decomposition, we extract frequency components from macroeconomic indicators in the Uncovered Interest Parity (UIP), Purchasing Power Parity (PPP), and Taylor Rule models. These components are then used as inputs for Lasso, Ridge, Elastic Net, and Random Forest regressions. Empirical results on GBP/USD, GBP/EUR, and GBP/CNY demonstrate three key findings: (1) frequency components significantly enhance exchange rate forecasting accuracy, with decomposed models generally outperforming their non-decomposed counterparts; (2) the selection of an optimal decomposition scale is crucial for improving predictive performance; and (3) the Taylor Rule-based models exhibit superior forecasting ability compared to UIP and PPP frameworks. Moreover, the proposed wavelet-based models outperform traditional benchmark methods, including the Random Walk and ARIMA, in most cases. These findings highlight the role of frequency decomposition in financial modelling and provide new insights into exchange rate predictability.
Keywords Exchange Rate Forecasting; Wavelet Decomposition; Economic Fundamentals Models; Machine Learning Regression; Time-Frequency Analysis

Dynamic Shrinkage and Selection for Exchange Rate Forecasting

Presenter: Yong Song (University of Melbourne)
Author(s): Yong Song;Ole Maneesoonthorn;Zheng Fan

WE propose a novel methodology that integrates dynamic sparsity and the dynamic shrinkage process into a coherent framework. This approach introduces an adaptive dynamic variable selection mechanism to achieve

time-varying sparsity, ensuring that variables contributing no meaningful variation are deactivated during specific periods. This feature enhances interpretability by highlighting key variables, removing noise, and maintaining model parsimony. The mechanism is incorporated into the dynamic shrinkage process, which balances global and local shrinkage on coefficient drift to better capture evolving coefficient patterns. By dynamically identifying relevant predictors based on their changing importance over time, this methodology adapts to shifting economic conditions and mitigates the potential risk of overfitting.

Signature Decomposition Method Applying to Pair Trading

Presenter: Jinghan Wang (Shandong University)
Author(s): Zihao Guo;Hanqing Jin;Jiaqi Kuang;Zhongmin Qian;Jinghan Wang

Quantitative trading strategies based on medium- and high-frequency data have long been of significant interest in the futures market. The advancement of statistical arbitrage and deep learning techniques has improved the ability of processing high-frequency data, but also reduced arbitrage opportunities for traditional methods, yielding strategies that are less interpretable and more unstable. Consequently, the pursuit of more stable and interpretable quantitative investment strategies remains a key objective for futures market participants. In this study, we propose a novel pairs trading strategy by leveraging the mathematical concept of path signature which serves as a feature representation of time series data. Specifically, the path signature is decomposed to create two new indicators: the path interactivity indicator segmented signature and the change direction indicator path difference product. These indicators serve as double filters in our strategy design. Using minute-level futures data, we demonstrate that our strategy significantly improves upon traditional pairs trading with increasing returns, reducing maximum drawdown, and enhancing the Sharpe ratio. The method we have proposed in the present work offers greater interpretability and robustness while ensuring a considerable rate of return, highlighting the potential of path signature techniques in financial trading applications.

Identifying Volatility Clusters, Common Trends and Comovement of DeFi Tokens

Presenter: Atikur Khan (North South University)
Author(s): Atikur Khan;Bashir Khan;Arifur Rahman

Investment decisions in the decentralized finance (DeFi) market are heavily influenced by the volatility dynamics of the cryptocurrency tokens and their comovement patterns. Understanding these patterns is crucial for investors seeking to optimize portfolios while managing risk in the face of high market fluctuations. This paper explores the volatility patterns and co-movement dynamics of DeFi tokens in cryptocurrency markets. Employing advanced time series models such as the generalized autoregressive conditional heteroscedasticity (GARCH) framework for estimating time-varying conditional variances, and clustering techniques with dynamic time warping and K-means clustering, this paper identifies distinct clusters of DeFi tokens with similar volatility profiles. Additionally, multichannel singular spectrum analysis (MSSA) is used to extract common volatility trends, filtering out noise and revealing co-movement dynamics within and across clusters. The results reveal distinct volatility clusters of DeFi tokens with varying degrees of intra-cluster correlation, highlighting the interdependencies and comovement among different token groups. The findings shed critical insights on the developed clustering and future trend identification scheme for future investment decision-making strategies, and a better understanding of volatility structures in the rapidly evolving cryptocurrency market.

Forecasting Extreme Values in Financial Time Series via Scores

Presenter: Ayla Jungbluth (Department of Mathematics, Ruhr-University Bochum)
Author(s): Ayla Jungbluth;Simon Trimborn;Johannes Lederer

Extreme value prediction in time series analysis remains a significant challenge for forecasting models, particularly in financial markets where tail events can have substantial impact. This paper introduces a novel time-dependent methodology specifically designed for forecasting extreme values in financial time series. Our approach addresses the limitations of conventional models by incorporating temporal dynamics into extreme value prediction. By combining

a time dependent Hüsler Reiss model with market informed parameter constraints, we keep the necessary parameters to estimate reasonable to handle the few observations in the tails of financial time series. We demonstrate the efficacy of our approach using an extensive dataset of individual stock components from the S&P 500 index. Our empirical analysis shows improvements in extreme value forecasting across diverse market sectors and varying conditions.

Evaluating Quarterly Inflation Forecasts in the EU: A Stepwise Analysis

Presenter: Katja Heinisch (katja.heinisch@iwh-halle.de)
Author(s): Katja Heinisch

This study provides a comprehensive evaluation of the inflation forecast performance of the European Commission (EC) across all European Union (EU) member states over the period 2000 to 2024. Forecasts are assessed over quarterly horizons extending up to eight quarters ahead and are benchmarked against Consensus Economics forecasts, where available. Accurate inflation forecasting remains essential for effective policy formulation, given its central role in guiding both monetary and fiscal decisions. The analysis examines the evolution of forecast accuracy over time and investigates whether performance differs systematically between euro area countries and Central and Eastern European Countries (CEECs). Findings reveal that EC forecasts tend to more accurately track inflation in the euro area, while inflation in CEECs has been more volatile and less predictable. Since 2021, inflation in CEECs has significantly exceeded that of the euro area. These developments emphasize the necessity of targeted national fiscal and structural policies, particularly in light of the limitations of the European Central Bank's (ECB) uniform monetary policy in addressing asymmetric shocks. The study further underscores the importance of data timeliness and frequency, particularly the use of monthly indicators, and explores the contribution of new nowcasting tools to improving short-term forecast accuracy and hence supporting economic resilience across the EU.

From Text to Forecast: A Large-Scale LLM Analysis of Central Bank Communication

Presenter: Thiago Silva (International Monetary Fund)
Author(s): Thiago Silva;Kei Moriya;Romain Veyrune

This paper investigates the role of central bank communication in forecasting inflation volatility. We develop a multilingual large language model (LLM) fine-tuned to classify central bank statements by topic, communication stance, sentiment, and audience. Applying this framework to a dataset of 66,854 documents from 157 central banks, we identify systematic shifts in communication following inflation-targeting adoption, with backward-looking discussions on exchange rates giving way to forward-looking statements on inflation and interest rates. Credibility tests confirm that these signals align with actual monetary policy stances. We construct a forward-lookingness score from monetary policy decisions and incorporate it alongside traditional macroeconomic variables in forecasting models of inflation volatility. We use a data-driven approach to select the best technique and the model hyperparameters by utilizing a horserace setup with rolling-window cross-validation schemes. While inflation persistence remains the dominant predictor, the forward-lookingness score is also relevant. SHAP-based interpretability analysis reveals that higher forward-lookingness scores predict lower subsequent inflation volatility, supporting theoretical models in which forward guidance enhances expectation anchoring and reduces uncertainty. Our findings provide empirical evidence that central bank communication serves as an effective monetary policy tool, actively contributing to inflation stability.

Forecasting Unemployment Rate: Phillips Curve, Services Prices, and Shelter Costs

Presenter: Dandan Liu (Kent State University)
Author(s): Dandan Liu

Unemployment rate is a key economic indicator and is closely watched by business, finance, and economic professionals. Accurate and timely unemployment forecasts assist business planners, financial institutions, and policymakers in anticipating economic trends. One widely used benchmark model for unemployment forecasting is the Phillips curve, which describes the tradeoff between inflation and economic activity. While the Phillips curve has faced criticism, it remains valuable for short-term modeling and forecasting. This study evaluates its out-of-sample

predictive power for U.S. unemployment by modifying the conventional model to incorporate key inflation components: service prices and shelter costs. The conventional Phillips curve model typically represents inflation using the Consumer Price Index (CPI), which includes both goods and services. However, prior research suggests that aggregate inflation may weaken the empirical estimates of the Phillips curve, as resource gaps tend to impact the costs of non-tradable services more than tradable goods. Studies have shown that distinguishing between service and goods inflation improves inflation modeling and forecasting accuracy. Building on this, the first modification in this study isolates the relationship between service inflation and the unemployment rate to enhance unemployment forecasting. The second modification examines the role of shelter costs, which constitute approximately 42% of the CPI. Shelter costs, including rent and owner-equivalent rent (OER), are highly correlated and respond uniquely to monetary policy shocks. Unlike most other CPI components, shelter costs tend to rise following contractionary monetary policy, potentially distorting the Phillips curve relationship. This study investigates whether excluding shelter costs from the inflation measure improves unemployment rate forecasts. A real-time forecasting exercise is conducted to compare the accuracy of different models, including the conventional and modified Phillips curve models, as well as other time-series forecasting approaches. The findings indicate that the modified models yield marginal improvements in forecasting accuracy across different time horizons.

Forecasting on the Global Methane Emissions Reduction by 2030 via Key Pathways in Inter-country Production and Consumption Networks

Presenter: Xiuli Liu (CHINESE ACADEMY OF SCIENCES, CHINA)
Author(s): Xiuli Liu;Zijie Cheng;Yuxing Dou;Mun Ho;Geoffrey Hewings;Shouyang Wang

Methane emissions embedded in global production and consumption follow complex, often obscured pathways, offering significant mitigation opportunities and posing substantial research challenges. To bridge this knowledge gap, we employed input-output and complex networks analyses to map annual methane emissions across global production-consumption networks by compiling extended Inter-country Input-Output (EICIO) tables, which covered 14 industries in 76 economies from 2000 to 2020. We developed an environmental double filtering method to identify the intricate network of connections among the leading embedded methane-emitting sectors within the (EICIO) tables, which represents the key pathways for embedded methane emissions. These pathways, just 0.3% of 1,132,096 total linkages, accounted for around 60% of global methane emissions. Our analysis revealed that emission intensity, per capita GDP, and technical coefficients significantly yet variably influenced these pathways' emissions. In our scenario analysis, we matched key pathways with feasible technologies to reduce methane emission intensity and assessed their associated costs, and identified specific linkages through which global methane emissions are predicted to be reduced by 30.7% from 2020 levels by 2030, with total costs estimated at US$20.63 billion. This study provides a practical strategy for developing effective and actionable plans to achieve the Global Methane Pledge by 2030 through international cooperation.

The Energy System Transition Pathway towards Carbon Reduction Using a Model-coupling Approach

Presenter: Jiali Zheng (Xi'an Jiaotong University)
Author(s): Jiali Zheng;Xiaoqing Hou;Jiaming Yang;Lianyang Jiao;D'Maris Coffman;Shouyang Wang

The energy system transition is widely regarded as an important strategy to achieve carbon reduction and is aligned with China's commitment to reach peak carbon emissions by 2030. Most modelling approaches in the existing literature do not pay sufficient attention to inter-sectoral dynamics. By using a model-coupling approach, this paper aims to forecast inter-sectoral energy consumption flows to 2030 and to explore energy system transition pathways at the national and city levels. The results show that historically heavy industries have consistently maintained a high share of energy consumption and emissions accounting for 49.9 % and 60.7 % respectively by 2021, mainly caused by direct energy-resource inputs rather than postprocessing inputs. In the scenario analyses, compared to the baseline scenario, the national EES scenario is predicted to reduce energy consumption by 6.7 % and emissions by 24.6 % in 2030, while the EES_CCS scenario can further reduce emissions by 48.4 %. Furthermore, the energy consumption and CO2 emissions across cities will be influenced by the industrial structure, the degree of electrification, and the amount of new energy installed.

Socio-Ecological Driving Mechanisms of Ecosystem Services and Territorial Spatial Ecological Function Optimization Strategies

Presenter: Gang Liu (Tianjin University)
Author(s): Gang Liu

Exploring the complex driving mechanisms of ecosystem services (ESs) is essential for achieving Sustainable Development Goals (SDGs). However, current research often ignores the influence of interactions between driving factors on ESs, leading to a major blind spot in ESs management. In this study, we take China as the research area to analyze the driving mechanism of ESs in 2005, 2010, 2015, and 2020. Firstly, we propose an algorithm to determine the optimal spatial scale of ESs driving mechanism research at the high-precision grid level. Then, we adopted the spatial econometric model and geographically weighted regression (GWR) model to analyze driving factors from the global and local perspectives. In addition, we use the Optimal Parameters-based Geographical Detector (OPGD) to analyze the interaction between driving factors and use the growth rate model to quantify the interaction intensity between driving factors. The results show that: (1) the driving factors of ESs show significant spatial spillover effects and spatial heterogeneity; (2) the interaction effect of driving factors is greater than that of single-factor explanatory power; (3) the interaction intensity between driving factors is different in different ESs. This study provides valuable insights into the driving mechanisms of ESs and informs the implementation of SDGs in China.

Diabetes-Specific Feature Integration for Stratified Risk Prediction and Interpretation of Acute Kidney Injury

Presenter: Siyu Dong (School of Public Health, Peking University)
Author(s): Siyu Dong;Xinwen Chen

Diabetic patients face a significantly elevated risk of acute kidney injury (AKI) due to chronic hyperglycemia and microvascular complications, making early risk stratification essential for clinical intervention. However, existing AKI prediction models often overlook diabetes-specific factors, focus only on binary outcomes, and are frequently limited by single-center data, constraining their generalizability. This study proposes a multicenter, stratified AKI risk prediction framework specifically designed for diabetic patients. Using 12,171 patients from the MIMIC-IV database, AKI stages were defined by KDIGO criteria. Both general variables (demographics, comorbidities, vital signs, laboratory tests) and diabetes-specific features (glycemic control, complications, antidiabetic medications, metabolic status) were integrated. LASSO regression was employed for feature selection, and a four-class prediction model was developed using Bayesian-optimized LightGBM. We further benchmarked the model against diverse machine learning algorithms under various resampling strategies, including SMOTE and NearMiss. The BO-LightGBM model outperformed other approaches, achieving a macro-AUC of 0.75, micro-AUC of 0.78, and an AUC of 0.89 for stage 3 AKI. SHAP analysis identified SOFA score, OASIS score, body weight, and antidiabetic medication use as key predictors, revealing complex nonlinear interactions. Moreover, feature importance varied across AKI stages: for instance, fasting glucose emerged as a critical predictor for stage 2 AKI. These findings provide mechanistic insights into the progression of AKI in diabetic patients and enhance the model's interpretability for clinical application. To ensure external validity, we conducted external validation on the eICU database using a minimal variable set, retraining the model on shared features. The external model achieved a macro-AUC of 0.66 and micro-AUC of 0.64, demonstrating moderate cross-database generalizability. This study confirms the potential of personalized prediction models for precision management of AKI risk in diabetic patients and provides a foundation for future clinical decision support systems and multicenter deployment.

Boosting Domain-specific Models with Shrinkage: An Application in Mortality Forecasting

Presenter: Li Li (北京科技大学)
Author(s): Li Li;Han Li;Anastasios Panagiotelis

This paper extends the technique of gradient boosting with a focus on using domainspecific models instead of trees. The domain of mortality forecasting is considered as an application. The two novel contributions are to use well-known stochastic mortality models as weak learners in gradient boosting rather than trees, and to include a penalty

that shrinks mortality forecasts in adjacent age groups and nearby geographical regions closer together. The proposed method demonstrates superior forecasting performance based on US male mortality data from 1969 to 2019. The proposed approach also enables us to interpret and visualize the results. The boosted model with age-based shrinkage yields the most accurate national-level mortality forecast. For state-level forecasts, spatial shrinkage provides further improvement in accuracy in addition to the benefits of age-based shrinkage. This improvement can be attributed to data sharing across states with large and small populations in adjacent regions and states with common risk factors.

Predicting the Future Impact of Weather Patterns on the Incidence of Infectious Diseases in Children

Presenter: Suryane Susanti (Faculty of Nursing, Universitas Indonesia)
Author(s): Suryane Susanti;Dessie Wanda;Holivia Jacinta;Arief Hakim;Atina Ahdika

Climate change is increasingly recognised as a significant factor affecting health, particularly in vulnerable populations such as children. As global temperatures rise and extreme weather events become more frequent, disaster risks such as flooding and drought also increase. This situation leads to more children being exposed to the impact of climate change as well as disasters. By examining historical data and current trends in weather patterns, this research aims to forecast the incidence of infectious diseases in children who are exposed to climate change-related disasters. We will accomplish this by constructing an integer-valued time series model for the number of cases of such diseases for several children's age groups. We will include the number of disasters as an explanatory variable potentially affecting these diseases. In particular, this research will focus on a case study in Indonesia, a tropical country that is highly vulnerable to climate change-related disasters and infectious diseases. The findings will provide insights into how future climate scenarios may increase health risks for children. This will help shape health strategies and policy responses, highlighting the need for proactive measures to protect future generations, particularly in Indonesia.

Decisions Reconciliation vs Forecasts Reconciliation

Presenter: Mahdi Abolghasemi (Queensland University of Technology)
Author(s): Mahdi Abolghasemi;Yingjie Zhao;Fotios Petropoulos

Hierarchical time series represent a set of time series that are naturally grouped into different levels of aggregation in a hierarchical structure. In practice, forecasts are often required for different levels of the hierarchy to be used for decision-making. Hierarchical forecasting generates coherent forecasts across the hierarchy, where forecasts at lower levels aggregate to forecasts at higher levels. In addition to forecast coherency, forecast reconciliation often improves average forecast accuracy (though accuracy might be compromised at certain hierarchy levels). Although reconciled forecasts are coherent across the hierarchy, the decisions resulting from such forecasts may not be. Furthermore, studies have shown that forecast accuracy might not always translate into better decisions quality. Hence, from a decision-making perspective, it may be worthwhile to focus on reconciling the decisions to be made instead. We propose to reconcile decisions and conduct an empirical analysis where forecasts (reconciled and unreconciled) are generated for M5 data. The inventory decisions that come from using these forecasts are then also reconciled, comparing their inventory performance with that of unreconciled inventory decisions. We present our findings from this empirical investigation and showcase the value of reconciled inventory decisions.

Probabilistic Nonlinear Forecast Reconciliation Methods: Projection, Importance Sampling, and Kalman Filtering

Presenter: Anubhab Biswas (University of Applied Sciences and Arts of Southern Switzerland (SUPSI))
Author(s): Anubhab Biswas;Lorenzo Nespoli;Giorgio Corani;Lorenzo Zambon

So far probabilistic forecast reconciliation has been studied assuming linear dependencies between time series. In this study, we focus on nonlinear reconciliation, in which the time series are linked by a nonlinear function f. We assume f to be known. We propose four probabilistic methods for nonlinear forecast reconciliation, which extend to the nonlinear case some methods already existing in the linear case. The Non-linear Bottom-Up (NL-BU) offers a simple baseline: it samples from the joint distribution of the independent (bottom) time series, and passes them through

the function f, providing a distribution for the dependent (upper) time series. The Non-linear Orthogonal Projection (NL-OP) method extends the linear projection framework to the nonlinear case by solving a constrained optimization problem using Newton-Raphson updates. We also adapt the Non-linear Bottom-Up Importance Sampling (NL-BUIS) approach for nonlinear dependencies, using the bottom-level distribution as a proposal and reweighting samples based on the upper-level likelihood. Finally, we propose a novel use of the Unscented Kalman Filter (UKF) to update bottom-level distributions using nonlinear information from the upper level, yielding coherent probabilistic forecasts. We evaluate all methods on a simulated case. We generate two time series via autoregressive processes and derive a third by summing their squared values. Results based on the energy score show that the UKF and NL-OP methods outperform both the base and bottom-up forecasts, achieving the highest probabilistic accuracy. In terms of log score at the bottom level, both UKF and NL-OP again perform best, with BUIS offering moderate improvements over the bottom-up approach. These findings highlight the UKF and nonlinear projections as strong candidates for probabilistic reconciliation in the nonlinear setting. Finally, the possibility and the relevance of using different probabilistic metrics to score forecasts for nonlinearly connected time series, based on the geodesic distance, is discussed.

Modeling the Uncertainty on the Covariance Matrix for Probabilistic Forecast Reconciliation

Presenter: Chiara Carrara (University of Pavia)
Author(s): Chiara Carrara;Lorenzo Zambon;Dario Azzimonti;Giorgio Corani

In forecast reconciliation, the covariance matrix of the base forecasts errors plays a crucial role. Typically, this matrix is estimated, and then treated as known. In contrast, we propose a Bayesian reconciliation model that explicitly accounts for the uncertainty in the covariance matrix. We choose an Inverse-Wishart prior, which leads to a t-Student reconciled predictive distribution and allows a completely analytical derivation. Empirical experiments demonstrate that this approach improves the accuracy of the prediction intervals with respect to MinT, leading to more reliable probabilistic forecasts.

Forecast Reconciliation with Heuristic Structure Averaging

Presenter: Jincen Luo (Beihang University)
Author(s): Jincen Luo;Bohan Zhang;Yanfei Kang

Forecast reconciliation has emerged as a key technique for ensuring hierarchical coherence in time series forecasting, particularly in domains such as tourism, energy, and macroeconomics. While existing forecast reconciliation methods have achieved promising results by ensuring coherence across hierarchical levels, these methods are typically limited to fixed hierarchical structures. Recent research has shown that by building upon traditional forecast reconciliation frameworks, constructing new hierarchical structures in a data-driven way provides a means to further enhance forecasting accuracy. In this study, we propose a novel reconciliation framework—forecast reconciliation with structure averaging—that systematically explores and synthesizes multiple hierarchical structures to improve forecast accuracy. Our method introduces a stochastic hierarchy jump process, guided by the Metropolis-Hastings algorithm, to sample from a space of regrouped hierarchies. Each structure yields its own reconciled forecasts, which are then aggregated via a weighted averaging scheme informed by structure-specific acceptance ratios. This approach enables the integration of structural uncertainty into the reconciliation process. Extensive experiments on both simulated datasets and real-world Australian labour force data demonstrate that our method consistently outperforms conventional base forecasts and reconciliation methods using the original hierarchy. These results underscore the value of data-driven structural exploration in hierarchical forecasting and offer a principled extension to existing reconciliation paradigms.

Risk spillover and Systemic Risk Early Warning in Financial Institutions: A DY-GNN-LSTM Model Based on High-Frequency Data

Presenter: Ke Zhang (Southwestern University of Finance and Economics)
Author(s): Ke Zhang;Haoran Liu

Accurate prediction and effective early warning of systemic risk are of paramount importance for maintaining

financial stability and promoting sustainable economic growth. This research aims to address the critical challenge of capturing the dynamic and interconnected nature of systemic risk by innovatively proposing a DY-GNN-LSTM hybrid model based on high-frequency financial data. Its primary objective is to dynamically predict risk contagion pathways among financial institutions and achieve early warning of systemic risk events.Compared to traditional methods, this model achieves significant progress through the seamless integration of three powerful techniques. First, the Diebold-Yılmaz (DY) spillover index is employed to quantify the dynamic interconnectedness and risk spillover effects among financial institutions. This provides a reliable measure of the strength and direction of risk transmission. Second, a graph neural network (GNN) is utilized to capture the complex network structure of the financial system, representing institutions as nodes and their interdependencies as edges. The GNN enables the model to learn the structural properties of the network and identify systemically important institutions. Finally, a long short-term memory network (LSTM) is incorporated to model the temporal dependencies and dynamic evolution of risk contagion. The LSTM captures the sequential patterns in risk transmission and enables the model to predict future risk levels.The DY-GNN-LSTM model operates in a sequential manner. First, the DY spillover index is calculated to quantify the risk interconnectedness among financial institutions. This information is then used to construct the financial network, which is fed into the GNN. The GNN learns the structural properties of the network and generates node embeddings that capture the risk characteristics of each institution. Finally, the LSTM uses these node embeddings, along with other relevant features, to predict future risk levels and identify potential systemic risk events.Empirical results demonstrate that the model can effectively identify key risk sources and risk contagion paths, significantly improving the accuracy and timeliness of systemic risk early warning.

Improved Random-Forest-Based Bilevel Optimization for Personalized Travel Demand Equilibrium under the E-R-A Framework

Presenter: Zixuan Zhao (Chongqing University)
Author(s): Zixuan Zhao;Kunhui Ye;Yuyao Liu

Under the coordinated evolution of high-quality urbanization and deep digital-intelligence integration, on-demand mobility services, exemplified by ride-hailing, have experienced rapid growth, enhancing travel immediacy and convenience and attracting many residents driven by personalized travel demand. However, the uneven spatiotemporal distribution of personalized travel demand increases the complexity of congestion management and results in a decline in service efficiency. Conventional strategies of traffic demand management have been ineffective in addressing above challenge. The research proposes a spatial equilibrium model for personalized travel demand distribution within an Efficiency-Reciprocity-Anisotropy (E-R-A) framework, employing complex network theory to capture the dynamics of a healthy and efficient traffic system. An improved Random-Forest-based Bilevel Optimization (RF-BO) algorithm is integrated to determine the critical equilibrium threshold. Empirical analysis is conducted using over 3.7 million ride-hailing records. The results indicate that personalized travel demand distribution exhibits both seasonal and stochastic fluctuations. The E-R-A model can determines the equilibrium threshold, beyond which traffic system efficiency, asymmetric interaction, and directional imbalance deteriorate significantly. This finding indicates why reducing overall demand alone is high-cost and ineffective in restoring service efficiency due to uneven demand distribution. Thus, a threshold-based management strategy is proposed that alleviates uneven demand distribution effectively enhances traffic system performance. This study provides a quantitative framework for understanding the impact of uneven personalized travel demand and offers actionable insights for improving service efficiency through targeted demand management.

Yield Nowcasting Reconcilition under Gaussian Framework

Presenter: Zhuoqun Xie (HUNAN UNIVERSITY)
Author(s): Zhuoqun Xie;Jin Yang;Ning Zhang;Daning Bi

Climate risk has become a crucial issue in the management of global agricultural risks. Doing a good job in “nowcasting” and early warning of meteorological disasters, as well as analyzing medium and long-term trends, accurately forecasting wheat yield, is of great significance for ensuring food security. Accurate wheat yield forecast is required at the regional, state, and county levels, and must take into account the impact of weather factors. However, the wheat yield from different counties and varieties are affected differently by weather, and predictions at different

levels may be inconsistent. Current research has paid less attention to the reconciliation and integration of global and regional yields of agricultural products. Therefore, under the condition of ensuring the consistency of forecasting, this paper provides an effective statistical forecasting model for the global forecast of multi-level wheat yield. Using spring and winter wheat from 1982 to 2022 in the United States as a case study, this paper discusses the use of point forecasting and probabilistic forecast reconciliation for multi-level wheat yield forecasting based on multidimensional weather factors. Scoring rules are used to make judgments on the forecasting consistency between levels. Finally, the value at risk is analyzed by forecasting wheat yield through probability reconciliation.

Policy Communication, Market Expectations, and Economic Resilience: Causal Identification Based on Mixed-Frequency Data and Textual Analysis

Presenter: Xinyu Guo (Jilin University)
Author(s): Xinyu Guo;Han Liu

Guiding the expectations of market participants through policies to enhance macroeconomic resilience has emerged as a crucial challenge in contemporary macroeconomic governance. Leveraging multidimensional data on policy communication, market expectation sentiment, and macroeconomic resilience, this study employs the mediation effect model and two advanced econometric models: the Mixed Frequency Data Sampling Time-Varying Parameter Local Projection Quantile Regression (MIDAS-TVP-LP-QR) and the Mixed Frequency Data Sampling Time-Varying Parameter Structural Decomposition Local Projection (MIDAS-TVP-SD-LP). By identifying the dynamic transmission mechanism from policy to expectation to resilience, we explore how government policy influences market expectations and enhances macroeconomic resilience through the channel of expected market sentiment. Our empirical results show the following: (1) policy significantly affects macroeconomic resilience through the expected sentiment of market participants, in which negative expectations have a stronger adverse impact on economic resilience in low resilience states, especially on cyclically sensitive sectors; (2) the sentiment tendency, subject clarity and operational directionality of the Policy Communication Index (PCI) play a key role in guiding the expected sentiment of the market, especially during economic downturns. Expansionary policies are more effective in boosting market confidence during economic downturns; (3) optimizing policy communication and expectation management can enhance the transmission effect of policy and improve macroeconomic resilience. The results of the study provide a valuable empirical foundation and policy insights for further improving the macroeconomic control tool system, optimizing the effect of policy communication and promoting high-quality economic development.

Can Public Opinion Big Data Help Improve Macroeconomic Forecasting?

Presenter: Xiaofen Li (School of Economics & Trade, Hunan University)
Author(s): Xiaofen Li;Hao Xiao;Junyi Yang;Xinjian Ye

This paper proposes a novel forecasting framework for predicting South Africa’s macroeconomy using big data on public opinion, based on deep learning methods. Using the GDELT database and incorporating data mining techniques and construction methods, effective information is extracted from 10.76 million news articles to derive international public opinion related to South Africa. The study integrates traditional macroeconomic indicators with public opinion big data, comparing the forecasting performance across econometric, machine learning, and deep learning models. Rolling out-of-sample prediction analysis indicates that the LSTM model achieves the highest forecasting accuracy. Subsequently, LSTM models with and without public opinion big data were developed, and the improvement in forecasting accuracy due to the inclusion of public opinion data was examined. This study demonstrates that public opinion big data enhances the accuracy of macroeconomic forecasts, particularly by improving the short-term forecasting capabilities of the model. Keywords: macroeconomic forecasting, public opinion, deep learning, big data

When Algorithms Speak with Accents: Economic Forecasting Disparities Across Language-Specific LLMs

Presenter: Mohan Xu (Peking University)
Author(s): Mohan Xu;Feng Li;Yao Tang;Yuankuan He

This paper investigates the transformative influence of AI-driven information acquisition and propagation via the Internet, with a particular emphasis on the growing prominence of large language models (LLMs) in digital information dissemination. While existing literature has yet to systematically address the intrinsic biases in LLM-generated narratives and their societal implications, our comparative analysis examines fundamental differences in responses to critical economic issues between two leading language-specific LLMs: the English-based ChatGPT and the Chinese-based Kimi. Our experimental design involves twice-daily queries using natural-language prompts during US and Chinese stock market closures, generating sequential economic forecasts encompassing macroeconomic outlooks, policy stances and bilateral relations between the two nations. Leveraging the comprehensive response datasets, we empirically assess the temporal responsiveness of LLMs to major macroeconomic shocks and examines systematic divergence in economic interpretation and forecasts between the two language-specific models when exposed to identical events. These findings elucidate mechanisms through which LLMs may induce cognitive disparities across linguistic user groups regarding macroeconomic trends, shedding light on how AI intermediaries might shape public discourse.

Crude Oil Futures Price Uncertainty Forecasting: An Explainable Deep Learning Framework Integrating News Sentiment

Presenter: Yaqi Mao (Nanjing University of Information Science and Technology)
Author(s): Yaqi Mao;Xiaobing Yu

Accurate forecasting of crude oil futures prices is essential for guiding energy investment decisions, managing market risk, and informing policy development. This paper proposes NewsSent GBO TFT, an interpretable deep learning framework that integrates artificial intelligence innovations with practical financial engineering applications. First, sentiment signals are quantitatively extracted from unstructured news text via a pretrained language model and fused with multidimensional economic variables—including supply–demand fundamentals, monetary and financial market conditions, and macroeconomic uncertainty measures—to construct a unified feature representation. Second, we present a collaborative feature selection strategy that leverages multiple evaluation criteria to assess feature importance from diverse angles, effectively eliminating redundant predictors and enhancing both model robustness and generalization. Finally, we employ a Gradient Based Optimizer (GBO) to jointly refine the architecture and hyperparameters of the Temporal Fusion Transformer (TFT), yielding precise point forecasts along with multi level confidence interval estimates. Empirical evidence indicates that NewsSent GBO TFT outperforms benchmark models in forecast accuracy and uncertainty quantification. Interpretability analyses show that short horizon forecasts are driven primarily by trading day effects and proximate temporal signals, while medium and long horizon forecasts increasingly incorporate macroeconomic dynamics and geopolitical risk factors. Overall, the proposed approach delivers transparent, high fidelity predictive insights that can significantly enhance decision making for investors, risk managers, and policymakers in volatile market environments.

Global and Regional Long-term Climate Forecasts: A Heterogeneous Future

Presenter: Lola Gadea Rivas (University of Zaragoza)
Author(s): Lola Gadea Rivas;Jesus Gonzalo

Climate is a long-term issue, and as such, climate forecasts should be designed with a long-term perspective. These forecasts are critical for crafting mitigation policies aimed at achieving one of the primary objectives of the Paris Climate Agreement (PCA) and for designing adaptation strategies to alleviate the adverse effects of climate change. Furthermore, they serve as indispensable tools for assessing climate risks and guiding the green transition effectively. This paper introduces a straightforward method for generating long-term temperature density forecasts using observational data, leveraging the realized quantile methodology developed by Gadea and Gonzalo (JoE, 2020). This methodology transforms unconditional quantiles into time series objects. The resulting forecasts complement those produced by physical climate models, which primarily focus on average temperature values. By contrast, our density forecasts capture broader distributional characteristics, including spatial disparities that are often obscured in mean-based projections. The proposed approach involves conducting an out-of-sample forecast model competition and integrating the forecasts from the resulting Pareto-superior models. This method reduces dependency

on any single forecast model, enhancing the robustness of the results. Additionally, recognizing climate change as a non-uniform phenomenon, our approach emphasizes the importance of analyzing climate data from a regional perspective, providing differentiated predictions to address the complexities of a heterogeneous future. This regional focus underscores the necessity of accounting for spatial disparities to better assess risks and develop effective policies for mitigation, adaptation, and compensation. Finally, this paper advocates that future climate agreements and policymakers should prioritize analyzing the entire temperature distribution rather than focusing solely on average values.

Improving Risk Forecast through Coverage Probability, Expectile-based Approach and Credibility Theory

Presenter: Khreshna Syuhada (Institut Teknologi Bandung)
Author(s): Khreshna Syuhada;Dessie Wanda;Syabil Syuhada

Forecasting future risk is an interesting and challenging topic in many fields particularly in energy finance, cryptocurrency market and healthcare market. In this study, we consider the problem of forecasting risk through risk measures such as Value-at-Risk (VaR) and Expected Shortfall (ES) and their assessment. The accuracy of such forecasts may be improved in several directions. First, we do taking into account the parameter variability so that we obtain an improved forecast with better coverage probability. Secondly, while most VaR and ES calculations are quantile-based, we carry out the expectile-based forecast due to the importance of magnitude loss. Third, we involve credibility theory of our forecast as we find out that risk profile may be treated as a random variable. We do some numerical analysis for energy and cryptocurrency markets as well as healthcare market to illustrate the improvement of risk forecast accuracy. Furthermore, we construct optimal portfolio of such markets and assess risk reduction of the resulting portfolios.

QRA-Based Ensemble Forecasting Method for Carbon Price Intervals

Presenter: Shunxin Ye (Xidian university)
Author(s): Shunxin Ye

Interval prediction possesses informational advantages by effectively forecasting and characterizing critical information such as carbon price fluctuation features and uncertainty. Ensemble forecasting can leverage the strengths of multiple models to mitigate risks including Perturbation Instability in Estimation (PIE) and the “putting all eggs in one basket” risk. However, naively integrating interval predictions faces challenges such as complex interval operations, cumulative error propagation, and error amplification. To address these issues, this study introduces Quantile Regression Averaging (QRA) into multiscale carbon price interval prediction research，and based on a decomposition-ensemble forecasting framework, we construct a QRA-based integrated carbon price interval forecasting methodology. The QRA integration method is simple and flexible, directly aggregating multiple independent point forecasts into interval predictions, thereby bridging the two forecasting paradigms. This approach not only further integrates the advantages of individual point prediction schemes but also avoids error accumulation, amplification, and subjective weight allocation issues inherent in traditional interval prediction integration processes . Additionally, it demonstrates robustness in addressing heteroscedasticity and leptokurtic/heavy-tailed distributions. The implementation strictly follows these steps: Decomposition: The original carbon price series is decomposed into simpler, stationary, and easily predictable modes using Variational Mode Decomposition (VMD).Mode-specific Forecasting: Models including ARIMA-GARCH, LSTM, NAR, LSSVM, and RVM are used to predict each mode. Model settings are adaptively optimized via Bayesian optimization algorithms driven by data characteristics.Point Forecast Integration: Predictions of all modes are aggregated to obtain point forecasts of the original carbon price.QRA-based Interval Generation: The five point forecasts are further integrated using QRA to generate carbon price interval predictions. Empirical studies based on European carbon market and Chinese carbon market data demonstrate that this method improves forecasting effectiveness. Compared with other models, it achieves higher prediction accuracy, effectiveness, and reliability.

Generalized Principal Component Analysis for Large-dimensional Matrix Factor Model

Presenter: Yujie Hou (Shandong University)
Author(s): Yujie Hou;Yong He;Haixia Liu;Yalin Wang

Matrix factor models have been growing popular dimension reduction tools for large-dimensional matrix time series. However, the heteroscedasticity of the idiosyncratic components has barely received any attention. Starting from the pseudo likelihood function, this paper introduces a Generalized Principal Component Analysis (GPCA) method for matrix factor model which takes the heteroscedasticity into account. Theoretically, we first derive the asymptotic distributions of the GPCA estimators by assuming the separable covariance matrices are known in advance. We then propose adaptive thresholding estimators for the separable covariance matrices and derive their convergence rates, which is of independent interest. We also show that this would not alter the asymptotic distributions of the GPCA estimators under certain regular sparsity conditions in the high-dimensional covariance matrix estimation literature. The GPCA estimators are shown to be more efficient than the state-of-the-art methods under certain heteroscedasticity conditions. Thorough numerical studies are conducted to demonstrate the superiority of our method over the existing approaches. Analysis of a financial portfolio dataset in the supplement illustrates the empirical usefulness of the proposed method.

Research on Volatility and Arbitrage Based on the AcGB2 Model

Presenter: Chen Jiayi (中国科学院大学)
Author(s): JIayi Chen

This paper proposes an improved volatility forecasting model—the Exponential Generalized Beta of the Second Kind GARCH (AcGB2) model—to address the nonlinear, asymmetric, and heavy-tail characteristics of asset prices and volatility in financial markets. While previous literature indicates the superiority of AcGB2 in actuarial and option pricing, theoretical developments have been limited to preliminary proofs and single dynamic structures. The proposed model introduces an external dual dynamic structure, offering enhanced capability in capturing fat tails, asymmetry, and volatility clustering. The theoretical part systematically analyzes the properties of the AcGB2 distribution, deriving the recursive structure of its dynamic location and scale under the GARCH framework, and constructing the conditional volatility equation. Through exponential transformation, the model better captures tail behaviors and stabilizes parameter estimation, while skewness and kurtosis control enhances adaptability to volatility clusters and extreme events. Parameters are optimized using Maximum Likelihood Estimation (MLE), confirming the model's strength in capturing volatility aggregation, heavy tails, and asymmetry. Empirically, the model is tested using simple stock returns to evaluate its in-sample and out-of-sample forecasting performance. Compared to traditional GARCH-type models (e.g., Student-t GARCH, EGARCH), the AcGB2 model demonstrates superior robustness, particularly in extreme market conditions. Additionally, it enables dynamic analysis of volatility spillovers and market risk premiums. Based on no-arbitrage pricing theory, the paper proposes arbitrage strategies including volatility arbitrage, cross-market arbitrage, and risk-spillover arbitrage, validating the model's applicability in complex market environments. This study expands the application of GB-class models in modeling asymmetry and fat tails, and innovatively uses simple return as the sample. It also provides a new methodology for extreme risk management and asset allocation optimization. In the future, combining the AcGB2 model with machine learning is expected to enhance its utility in analyzing cross-asset volatility transmission and real-time risk prediction.

Dynamic Factor Correlation Model

Presenter: Chen Tong (Xiamen University)
Author(s): Chen Tong;Peter Hansen

We introduce a new dynamic factor correlation model with a novel variation-free parametrization of factor loadings. The model is applicable to high dimensions and can accommodate time-varying correlations, heterogeneous heavy-tailed distributions, and dependent idiosyncratic shocks, such as those observed in returns on stocks in the same subindustry. We apply the model to a "small universe" with 12 asset returns and to a "large universe" with 323 asset returns. The former facilitates a comprehensive empirical analysis and comparisons and the latter demonstrates the flexibility and scalability of the model.

Time Series Forecasting Using Quantile Regression Based on Fuzzy-probabilistic Inference

Presenter: Tomas Tichy (VSB-TUO)
Author(s): Tomas Tichy;David Neděla;Michal Holčapek;Nhung Cao;Radek Valasek

Weighted quantiles are essential tools in statistical analysis and regression, particularly for time-series data and moving quantile functions. Traditionally, their computation has relied on linear programming methods, which, while effective, can be computationally intensive. In this paper, we introduce an alternative approach that utilizes the right derivatives of the associated piecewise linear function. By minimizing this function, we derive the corresponding weighted quantile. This alternative approach retains the precision and reliability of traditional techniques while simplifying the computation. Furthermore, we combine this approach with probabilistic fuzzy rules and inverse quantile fuzzy transforms to propose a framework for estimating running quantiles and predictive modelling. The effectiveness of the method is demonstrated through both algorithmic implementation and empirical analysis using various financial data.

Modeling a Virtual Reality Enabler: a Climate Change Use Case

Presenter: Anna Boros (University of Warsaw)
Author(s): Anna Boros;Szymon Talaga;Mikołaj Biesaga;Andrzej Nowak

Virtual Reality (VR) offers an innovative approach to addressing climate change and environmental degradation through gamification and user engagement. Our research focuses on developing Guest XR, a VR agent designed to promote sustainable behavior within the context of a common-pool resource (CPR) dilemma. In CPR scenarios, non-excludable but rivalrous resources are overused, leading to environmental degradation. By integrating VR, we provide a "learning-by-doing" experience where players directly interact with the problem. To achieve this, we first create an agent-based simulation to model CPR dynamics in a virtual environment. From existing literature, we identify three key factors that influence group behavior: (i) agents' time perspective, (ii) agents' perception of environmental conditions, and (iii) their social orientation. Guest XR is designed to influence these factors through adaptive interventions, guiding players toward sustainable actions. A key feature of this approach is the forecasting function of each agents that GUest XR can influence by changing the level of three parameters described above. This function enables agents to forecast the amount of resources they should harvest at each turn of the game. By incorporating a forecasting model, agents can predict future resource availability based on current usage patterns, helping players make informed decisions about resource consumption. This feature encourages players to balance short-term gains with long-term sustainability, fostering cooperation in managing shared resources. The results of our CPR simulations are promising. The virtual game mirrors real-world behavioral patterns observed in field experiments, including the emergence of unsustainable cycles and successful resource management strategies. This demonstrates that our VR environment, governed by Guest XR, effectively mirrors the complex dynamics of real-world resource dilemmas. Moreover, by integrating the forecasting function, players are empowered to make more strategic decisions, leading to improved cooperation and better resource management outcomes.

Understanding Spatio-Temporal Extremes: Max-Stable Models in Climate Change Research

Presenter: Caston Sigauke (University of Venda)
Author(s): Caston Sigauke

The study addresses the increasing frequency and intensity of climate and weather extremes, particularly drought risk, due to high temperatures and low precipitation. Through the use of max-stable models in analyzing spatio-temporal extremes, the study applied Hopkin's statistic to determine whether clustering exists in the data. It presented a high degree of clusterability, with Hopkin's values of 0.7317 for rainfall and 0.8446 for temperature. Hierarchical clustering delineated the study region into three temperature and precipitation clusters. Several max-stable process models were subsequently fitted across them, among which the Schlather model with multiple covariance functions was better than the Smith model, particularly its Gaussian covariance function. The findings aim to enhance drought dynamics knowledge triggered by extreme temperatures and light rainfall, providing a helpful output for hydrologists, meteorologists, climatologists, and agricultural policymakers. In addition, such a modelling approach could be utilized to evaluate climate change's effect on hydrologic extremes. Throughout the study, the importance of accounting for

the spatial size and geometry of regions during the extremal dependence calculation for gaining insight into the spatial evolution of climate-related risks prevalent across different locations.

Prediction of Children′ s Health Status when Impacted by Climate Change

Presenter: Dessie Wanda (Faculty of Nursing, Universitas Indonesia)
Author(s): Dessie Wanda;Suryane Susanti;Holivia Jacinta;Prihatini Novitasari

Climate change can cause detrimental effects on children's health in the form of direct or indirect effects. We investigate the potential of the severity of children's health status when impacted by climate change through prediction by using logistic regression algorithm. It incorporates various factors, including the type of climate change, the children's living conditions, the duration of their residency in a particular location, parental concerns about climate change, and parental actions taken to mitigate its impact. We anticipate that the algorithm will provide accurate predictions of health outcomes and identify key risk factors that contribute to adverse health effects in children. This insight can guide policymakers and healthcare providers in developing targeted interventions to protect vulnerable populations from the impacts of climate change.

On the Estimation of Climate Normals and Anomalies

Presenter: Tommaso Proietti (University of Rome Tor Vergata)
Author(s): Tommaso Proietti;Alessandro Giovannelli

The quantification of the interannual component of variability in climatological time seRies is essential for the assessment and prediction of the El Ni˜no - Southern Oscillation phenomenon. This is achieved by estimating the deviation of a climate variable (e.g., temperature, pressure, precipitation, or wind strength) from its normal conditions, defined by its baseline level and seasonal patterns. Climate normals are currently estimated by simple arithmetic averages calculated over the most recent 30-year period ending in a year divisible by 10. The suitability of the standard methodology has been questioned in the context of a changing climate, characterized by nonstationary conditions. The literature has focused on the choice of the bandwidth and the ability to account for trends induced by climate change. The paper contributes to the literature by proposing a regularized real time filter based on local trigonometric regression, optimizing the estimation bias-variance trade-off in the presence of climate change, and by introducing a class of seasonal kernels enhancing the localization of the estimates of climate normals. Application to sea surface temperature series in the Ni˜no 3.4 region and zonal and trade winds strength in the equatorial and tropical Pacific region, illustrates the relevance of our proposal.

Dynamic Options Portfolio Construction with Very Short Time-to-Maturity

Presenter: Maciej Wysocki (University of Warsaw)
Author(s): Maciej Wysocki

This paper examines the construction and performance of portfolios with a very short time to expiration, where positions in options are adjusted daily to maintain a fixed time to expiration. It also explores various methods for determining the number of contracts in the portfolio, including Kelly criterion and VIX-based sizing (see Cboe Global Markets, 2018). Additionally, the study presents an approach to hedging the portfolios using the Black-Scholes-Merton model (Black and Scholes, 1973; Merton, 1973). The primary objective of this study is to develop and empirically evaluate a novel approach to options trading that aims to optimize portfolio performance by maintaining a consistent time-to-maturity for the options in the portfolio. This strategy involves daily rebalancing of options positions to ensure a fixed expiration date, presumably allowing for more predictable and controlled risk management. Utilizing options near expiration allows harvesting theta and volatility-risk premiums. These strategies are compared against holding a benchmark index to assess the effectiveness of the new approach under various market conditions. Using 1-minute S&P 500 Index options data from 2018 to 2023, we benchmark hedged and unhedged option-writing strategies against buy-and-hold, considering risk-adjusted performance metrics and transaction costs. The results indicate that it is possible to construct strategies that outperform selected benchmarks in terms of risk-adjusted performance

metrics. Our strategies reveal the great potential of short-dated index options and showcase the high-profit potential of option-writing strategies.

Enhancing volatility prediction with the Leave-one-out kernel density estimates for outlier detection in GARCH models

Presenter: Wei Chean Tan (Universiti Putra Malaysia)
Author(s): Wei Chean Tan;Wei Chong Choo;Le Li

Many research has presented the existing of outliers in time series data causing errors and biases in model optimisation. This will lead to the decrement of the prediction accuracy for volatility forecasts. Hence, outlier detection is important in identifying the outlier for minimise the adverse effects from outlier. The new outlier detection approach, lookout was introduced which this approach detects outliers using leave-one-out kernel density estimates and extreme value theory. Furthermore, lookout identify the bandwidth suitable for outlier detection without any user input by implements persistent homology. Lookout is combining with GARCH, GJR-GARCH, IGARCH, STES-IGARCH, GARCH-MIDAS by identify the outlier in time series data before input into the models for optimisation. The research employs evaluation metrics of mean absolute error (MAE) and root mean square error (RMSE) to evaluate the robustness and performance consistency model. Finding presents the predictive capacity of the models is influenced by the outlier. The results are showing that the prediction performance for models GARCH, GJR-GARCH, IGARCH and STES-GARCH were improved by implementing the outlier detection approach but GARCH-MIDAS yet to show improvement on this. The prediction performance for GARCH-MIDAS is outperform GARCH and GJR-GARCH.

Using Quantile Time Series and Historical Simulation to Forecast Financial Risk Multiple Steps Ahead

Presenter: Richard Gerlach (The University of Sydney)
Author(s): Richard Gerlach;Giuseppe Storti;Antonio Naimoli

A method for quantile-based, semi-parametric historical simulation estimation of multiple step ahead Value-at-Risk (VaR) and Expected Shortfall (ES) models is developed. The method is primarily based on employing the quantile loss function, analogous to how the quasi-likelihood is employed by standard historical simulation methods. The estimated quantile series is used to scale the returns data, then re-sampling is employed to estimate the forecast distribution one and multiple steps ahead, allowing tail risk forecasting. The method is extended to allow a measurement equation, thus incorporating realized measures and including Realized GARCH and Realized CAViaR type models in the class of models it pertains to. The proposed method implicitly assumes, and is applicable to, any data or model where the relationship between VaR and ES in the conditional return distribution does not change over time; this includes most modern financial time series models. The finite sample properties of this method, and its comparison with existing historical simulation methods, are evaluated via a simulation study. A forecasting study, applied to 3 indices and 3 assets, assesses the relative accuracy of the 1% and 2.5% VaR and ES one-day-ahead and ten-day-ahead forecasting results for the proposed class of models compared to several competitors.

China's GDP Nowcast and Business Cycle Analysis - Based on the Three-regime Markov Switching Dynamic Factor Model

Presenter: Xiao Cai (University of Chinese Academy of Social Sciences)
Author(s): Xiao Cai;Liang Jin

In the context of increasingly abundant data, economic nowcasting research has become more frequent and immediate. Quarterly released GDP data no longer satisfies the needs for research and prediction using quarterly data. This paper uses Elastic Net to reduce dimensions after preliminary screening of hard indicators and survey indicators, incorporating the twice-screened parameters into the three-regime Markov switching dynamic factor model, and estimating various parameters. Compared to traditional two-regime models, the three-regime model decomposes economic downturns into general downturns and severe downturns. The estimation results show that the three-regime model can identify severe economic downturns and predict economic turning points in advance, thus providing

early warnings for economic crises. Additionally, regardless of whether the out-of-sample period includes pandemic times, the two-regime and three-regime Markov regime switching dynamic factor models (especially the three-regime model) exhibit significantly superior real-time and short-term forecasting capabilities compared to the benchmark AR(1) model, providing more accurate predictions for GDP growth rates.

Forecast Reconciliation and Hierarchical Structures in Subnational Mortality Forecasting: Insights from Japan

Presenter: Chenlu Deng (College of Finance and Statistics, Hunan University)
Author(s): Chenlu Deng;Lingyu He;Ning Zhang

Accurate mortality forecasting at national and subnational levels is crucial for public policy planning, resource allocation, and actuarial assessments. Ensuring coherence across forecasts at different hierarchical levels is essential in this process. While forecast reconciliation techniques have been widely applied in this field, not all advanced reconciliation methods have been tested on subnational mortality datasets, leaving their effectiveness in such contexts uncertain. Moreover, existing studies largely rely on predefined hierarchical structures, with limited exploration of how structural variations impact forecasting performance. This study focuses on subnational age-specific mortality forecasting in Japan, introducing the empirical minimum trace (EMinT) reconciliation method to this dataset for the first time. We systematically compare EMinT with other widely used reconciliation techniques, comprehensively evaluating their forecasting performance while also examining the influence of hierarchical complexity on reconciliation effectiveness. The results indicate that at the national level, the EMinT method consistently outperforms other approaches, with its forecast accuracy further improving as the hierarchical complexity increases. However, when both national and regional levels are considered simultaneously, the bottom-up approach demonstrates overall superior performance. This study validates the applicability of the EMinT method in subnational mortality forecasting for Japan and reveals the dynamic relationship between hierarchical complexity and forecast accuracy, providing valuable insights for demography, actuarial analysis, and policy development.

Beyond Numbers: Enhancing Crude Oil Price Forecasting with Economic and Sentiment Signals via LSTM-Attention Networks

Presenter: Shi Yuanyuan (Research Center of Management Science and Engineering, Jiangxi Normal University)
Author(s): Yuanyuan Shi;Mingchen Li

Abstract: Crude oil price forecasting is vital for economic decision-making but remains challenging due to its nonlinear, volatile, and multifactorial nature. This study proposes a hybrid deep learning framework that integrates macroeconomic variables with news-based sentiment signals to enhance predictive performance. A Long Short-Term Memory (LSTM) network with a self-adaptive attention mechanism is employed to model temporal dependencies and dynamically highlight influential time steps. Sentiment features are extracted from Google News using a fine-tuned CrudeBERT model, following standard text preprocessing procedures. Experimental results on Brent and WTI datasets show that the proposed model significantly outperforms traditional time series and machine learning approaches. The attention mechanism improves both accuracy and interpretability by revealing the temporal and feature-level importance in prediction. This research highlights the benefits of combining structured and unstructured information within an adaptive deep learning architecture for more accurate and transparent crude oil price forecasting.
Keywords: Crude oil price forecasting; LSTM-attention network; Sentiment analysis; Deep learning

A Drift-Aware Dynamic Ensemble Learning Framework for Forecasting Agricultural Futures Returns

Presenter: Zeng Liling (School of Computer Sciences, Guangdong Polytechnic Normal University)
Author(s): Zeng Liling;Zhang Dabin

Agricultural futures prices possess crucial functions of hedging and price discovery, analyzing their fluctuation trend and making scientific prediction can provide accurate and efficient information service for agricultural production and circulation, and promote the healthy and sustainable development of agricultural industry. However,

the complexity and time-varying nature of agricultural futures markets pose significant challenges for forecasting. Although ensemble models have been proven to enhance prediction reliability, addressing the relative performance variations among models and shifts in data distribution to generate effective prediction combinations remains a critical research issue. This study proposes a drift-aware ensemble learning framework that dynamically selects the most effective models and performs time-varying integration at each time step. First, thirty candidate models are generated by integrating data processing techniques with multiple forecast models to comprehensively capture sample information. Second, an initial selection process of candidate models is dynamically executed utilizing a performance drift detection mechanism. Following each drift detection, a second-stage selection is performed given the significance of diversity in ensemble models. Finally, final predictions are calculated by combining the outputs of selected models via a sliding-window weighted average. The effectiveness of this framework is evaluated using four sets of Chinese agricultural futures return index data. In the prediction of corn, soybean meal, cotton and rubber futures, the MAPE error is 0.1642%, 0.3965%, 0.2471% and 0.3515% respectively, which belongs to high-precision prediction. Compared with simple mean integration, the error reduction rates of MAPE are 73.11%, 94.50%, 69.98% and 94.68%, respectively. Empirical results demonstrate that the proposed model exhibits significant performance advantages in both level accuracy and directional accuracy.

Simultaneous Variable Selection and Estimation for a Partially Linear Cox Model

Presenter: Tingting Cai (Capital Normal University)
Author(s): Tingting Cai;Mengqi Xie;Tao Hu;Jianguo Sun

We consider simultaneous variable selection and estimation for a deep neural network-based partially linear Cox model and propose a novel penalized approach. In particular, a two-step iterative algorithm is developed with the use of the minimum information criterion (MIC) to ensure sparse estimation. The proposed method circumvents the curse of dimensionality while facilitating the interpretability of linear covariate effects on survival, and the algorithm greatly reduces the computational burden by avoiding the need to select the optimal tuning parameters that is usually required by many other popular penalties. The convergence rate and asymptotic properties of the resulting estimator are established along with the consistency of variable selection. The performance of the procedure is demonstrated through extensive simulation studies and an application to a myeloma dataset.

A Multi-Crop Yield Prediction Framework Integrating Explainable Ensemble Learning and Heterogeneous Data Sources

Presenter: Boting Zhang (College of Mathematics and Informatics, South China Agricultural University)
Author(s): Boting Zhang;Dabin Zhang;Liling Zeng;Huanling Hu

Accurate prediction of crop yields assumes a crucial role in scientifically guiding agricultural production activities and maintaining the stability of the global food supply. To address the limitations of existing single prediction models in adapting to the biological characteristic variations among different crops and capturing nonlinear interactions between crops and multiple environmental factors, this study developed a multi-crop universal prediction framework based on Stacking ensemble learning. The framework integrates multi-source heterogeneous data encompassing 11 meteorological parameters, 8 soil attributes, and 2 remote sensing indices, totaling 629 features. The Pearson correlation coefficient was employed for feature selection to eliminate redundant and noisy features. By aggregating prediction results from ten heterogeneous models (including statistical models and machine learning models), the framework achieved dynamic adaptation to crop-specific growth characteristics and precise yield prediction. Experimental results on the county-level crop yield dataset in the United States reveal that the proposed framework has exhibited remarkable prediction performance in the prediction tasks of corn, soybeans, and winter wheat, indicating the robustness of the proposed framework. Furthermore, to systematically investigate the synergistic effects of multiple environmental factors on crop yield, this study implemented the Shapley Additive Explanations (SHAP) method. Quantitative evaluations were conducted from global and local interpretability perspectives to assess each feature's impact on yield. The SHAP analysis reveals that meteorological factors exert a more significant influence on the yield, particularly precipitation and average temperature. These findings align with established crop growth mechanisms, thereby validating the model's interpretability.

Inter-Provincial Social Insurance Revenue and Expenditure Forecast in China

Presenter: Di Zheng (Beihang University)
Author(s): Di Zheng;Haitao Zheng;MIngzhe Zhang

China's social insurance system has not yet achieved nationwide pooling, and significant disparities exist among provinces in terms of current and accumulated fund balances. Coupled with differences in demographic structures, economic development levels, and social insurance systems across provinces, the future evolution of social insurance revenues and expenditures is expected to exhibit pronounced regional heterogeneity. This paper constructs a multi-regional population forecasting model that incorporates both inter-provincial migration and intra-provincial urban–rural migration, to systematically project the demographic evolution of all provinces in China. Based on the population forecasts, a social insurance actuarial model is employed to dynamically forecast the future social insurance revenue and expenditure status of each province. The key findings are as follows: (1) Most provinces will experience a continuous decline in total population over the coming decades, with the three northeastern provinces (Heilongjiang, Jilin, and Liaoning) experiencing the most significant declines due to low fertility rates and net out-migration; (2) Population aging will deepen across all provinces, with the aging rate expected to remain between 40% and 45% in most provinces by 2100, and particularly severe in Beijing, Tianjin, and the three northeastern provinces; (3) As a result of population aging, the social insurance revenue-expenditure gap in all provinces will continue to widen, with the gap exceeding 10% of regional GDP in some provinces such as Beijing, Shanghai, Henan, and Hainan; (4) There are marked differences in the solvency capacity of social insurance funds across provinces—provinces like Shandong and Fujian are expected to exhaust their cumulative balances within the next five years, facing imminent payment crises, while western provinces such as Yunnan, Tibet, Gansu, Qinghai, and Xinjiang maintain relatively sufficient cumulative balances to cover their fund gaps for the next 20 years. The findings of this paper provide quantitative support and policy insights for advancing nationwide pooling and optimizing the inter-regional social insurance adjustment mechanism. Keywords: Population migration; Multi-regional population forecasting; Social insurance revenue and expenditure; Insurance actuarial

Nowcasting Tail Risk of GDP Growth under High-dimensional Data: A Novel Approach Combining MF-QRLSTM with a Weighted Composite Probability Distribution

Presenter: Han Liu (Jilin University)
Author(s): Han Liu;Lijun Wang

Nowcasting tail risk has become a critical tool in modern economic analysis, particularly in the context of increasing economic uncertainty and volatility. Tail risk refers to the probability of extreme negative outcomes in economic indicators, such as sharp declines in GDP growth crashes. Traditional forecasting methods often struggle to capture these rare but impactful events due to their reliance on low-frequency data and delayed reporting. Nowcasting, by contrast, leverages high-frequency data to provide real-time estimates of economic conditions and potential risks. This approach is especially valuable for policymakers, investors, and financial institutions who need timely insights to mitigate risks and make informed decisions. We draw on Adrian et.al (2019) definition of economic tail risk but address two key limitations. First, traditional quantile regression models are constrained by their linear form and struggle to handle high-dimensional mixed-frequency data. To overcome this, we propose an MF-QRLSTM model to fit the quantiles of the predictive distribution. Second, given the inherent uncertainty in economic growth distributions, assuming that economic development follows a specific distribution is often unrealistic. Therefore, we introduce a composite probability distribution with weighting coefficients to fit the quantiles of the economic growth distribution. This paper focuses on nowcasts of tail risk to China's GDP growth, with a potentially wide array of monthly information used to produce nowcasts on a quarterly basis. Our results show that the weighted composite probability distribution forecasting model exhibits superior prediction accuracy and risk characterization capabilities across different forecast horizons. Within certain limits, additional information enhances the accuracy of nowcasts of tail risk in GDP growth. Accuracy generally improves as more data becomes available within a quarter, with monthly data playing a more critical role than quarterly data. Additionally, we observe that industrial output, consumption, and expectation factors significantly contribute to the prediction of economic tail risk. These findings provide critical references for real-time monitoring tail risk of GDP growth.

Forecasting Demand Surges for Irregular Events in Ride-Hailing Services Using Leading Indicators

Presenter: Marc Schwalbach (Grab)
Author(s): Marc Schwalbach;Ruike Zhang;Zhifei Shao

This study explores the potential of leading indicators, including event information, scheduled bookings, price checks, and pre-event drop-offs, to enhance demand forecasting for irregular events (e.g., concerts, sports matches) in the ride-hailing sector. These events present significant forecasting challenges due to their varying degrees of popularity, lack of historical data, and limited seasonality. We demonstrate that including these leading indicators leads to significant forecasting improvements compared to adding average historical demand uplift to the forecast results of established benchmarks, such as Seasonal Naive, AutoARIMA, and Prophet. We consider adding the leading indicators as exogenous variables, where applicable, as well as deriving an additive uplift factor that is added to the base forecast results. This improvement can be further distinguished by various time intervals between the creation of leading indicators and the conclusion of events, noting some leading indicators may only be available closer to the actual event. These findings can be broadly applied to ride-hailing demand forecasting use cases, including supply planning and anomaly detection, as well as for identifying missing data in event data sources.

I Don’t Understand Your Model; Explain it to Me!

Presenter: Trevor Sidery (Tesco Plc)
Author(s): Trevor Sidery

To bring the benefit of automated forecasts to a business it is important that everyone trusts the outcome of any forecasting solution. One key part of this to have explainable models. There has been a lot of discussion in the forecasting community about how to make sure that models are explainable, but these are often from a mathematical position, and not always the needs of the business. Similarly, unless business stakeholders understand the trade-offs that forecasters will need to make to fulfill various requests to improve explainability, we may end up with a non-ideal solution. I will try to shed some light on all these aspects using some examples we have come across over the last few years

Forecasting-driven Anomaly Detection at Scale: Lyft’s Approach to Observability

Presenter: Han Wang (Lyft & Nixtla)
Author(s): Han Wang;Anindya Saha

At Lyft, a leading ride share app, machine learning models power millions of high-stakes decisions daily, influencing physical safety, fraud detection, pricing optimization, and operational efficiency. Consequently, ensuring model reliability is crucial for business performance. Monitoring such diverse and complex ML systems at scale presents unique challenges, including distributional shifts in model inputs, unexpected feature definition changes, upstream bugs, or other unexpected events. To tackle these challenges systematically, we developed a forecasting-driven observability and anomaly detection solution. First, we profiled heterogeneous model features and outputs into standardized time series, significantly reducing the problem space and allowing us to rapidly detect deviations in metrics like call volumes, prediction scores, and input distributions. Furthermore, this standardization enables rigorous evaluation of diverse forecasting methodologies, ranging from classical statistical approaches (e.g., ARIMA, ETS, MSTL) to advanced neural forecasting methods (e.g., DeepAR, Transformers, NHits), as well as recent innovations involving state-of-the-art pre-trained foundation models. As a key result of this methodology we achieved a substantial reduction in false positives, leading to improved efficiency and trust in anomaly alerts. In this paper, we will deep dive into the methodological and practical lessons that we gained while developing and deploying this time series system into production, ensuring comprehensive coverage across models, reducing detection latency and human effort, enabling rapid onboarding of new ML models, and even including key business metrics. Additionally, this approach enables real-time anomaly detection, which further facilitates immediate anomaly alerts, prompting rapid root-cause investigation and remediation. This paper will illustrate Lyft’s experience in creating a state-of-the-art forecasting-driven anomaly detection pipeline, offering practitioners concrete strategies for building unified, scalable

observability solutions to reliably maintain model quality in complex production settings.

Candlestick Components and Market Directions: Not All Components Are Created Equal

Presenter: Haibin Xie (University of International Business and Economics)
Author(s): Haibin Xie;Kenan Qiao

This paper investigates if candlestick components are informative for predicting U.S stock market directions. We obtain three interesting findings: first, candlestick components are informative only when they are combined with market trend; second, not all candlestick components are created equal. It is the lower shadow which significantly and negatively predicts the market directions; third, it is the unexpected lower shadow which is informative for predicting market directions. We also explore the economic value of lower shadow for predicting market directions, and the results demonstrate significant economic gains of lower shadow forecasts relative to the simple buy and hold investment strategy.

A Flight-Level Neural Network Approach to Forecasting Daily Airport Fuel Uplift

Presenter: Alex Cao (United Airlines)
Author(s): Alex Cao;Andrew Alvarez;Amy Geraghty;Shawnae Johnson;Rich Alster

How many gallons of fuel will United Airlines put, or uplift, onto its flights departing an airport on a future date is a crucial operational forecast for inventory control. Previously, ARIMA-based time series modeling of daily airport aggregate uplift was sufficiently accurate. As the airline's business needs evolve to be more dynamic, we're interested in forecasting more accurately and on a longer time horizon. To that end, our approach has shifted drastically to forecasting individual flight uplifts to successfully capture exogenous variables like wind, weight, economic tankering opportunity, and remaining fuel at the effective granularity. By utilizing flight features derived from historical patterns of wind, passenger booking data, origin and destination fuel prices, as well as the same for the inbound flight, we more accurately forecast daily airport aggregate uplift (by summing over flights). In addition, we circumvent the recency assumptions and biases of time series models and our forecast horizon is less constrained. Our flight-level neural network approach reduces daily airport uplift error by 15.7% on average, compared to the daily airport aggregate uplift time series model.

Have Arrivals and Departures in Mexico Returned to Pre-pandemic Normality: the Case of Mexico City International Airport?

Presenter: Erick Heredia (Universidad Anáhuac México)
Author(s): Erick Martin Heredia Anaya;Jose Eliud Silva Urrutia

This study analyzes the recovery of passenger flows at Mexico City International Airport (AICM), one of the busiest airports in Latin America, following the effects of the COVID-19 pandemic. The objective was to determine whether domestic and international arrivals and departures at the airport's two terminals had returned to pre-pandemic levels. Two methodological approaches were applied: a vector autoregressive model (VAR) using monthly data from March 2012 to December 2019, and a hierarchical model using data from February 2015 to December 2019, both validated using statistical tests such as T-test, Breusch-Pagan, Ljung-Box and Shapiro-Wilk. The results indicate an uneven recovery; while Terminal 1 shows significant setbacks in domestic flights, with figures similar to those of more than a decade ago, and stagnation in international flights, Terminal 2 shows a more robust recovery in domestic flights, albeit with limited international growth. Finally, using the Diebold-Mariano test, it is concluded that the VAR model is statistically superior to the hierarchical model, suggesting the need for differentiated strategies to promote the full recovery of air traffic.

A Spatio-Temporal Dynamic Graph Model Based on Dual-Stream Convolutional Neural Network for Traffic Speed Prediction

Presenter: Xudong Long (South China Agricultural University)
Author(s): Long xudong

Accurate traffic speed prediction is crucial for the development of intelligent transportation systems in urban areas. However, the task of precise vehicle speed prediction is highly challenging due to the intricate spatio-temporal dependencies inherent in traffic data. Firstly, while some studies have attempted to incorporate additional factors influencing traffic flow speed, such as flow and weather conditions, into their models, the underlying mechanisms remain inadequately defined. This lack of clarity hampers the ability to capture the latent patterns within traffic flow. Secondly, the nonlinearity and high-dimensional nature of flow speed data, coupled with the dynamic characteristics of traffic networks, often limit the predictive accuracy of individual models. To address these challenges, this study introduces a novel, comprehensive, and highly adaptive hybrid deep learning framework named DSAG-GCGRU. Specifically, we employ a dual-channel convolutional neural network (CNN) to investigate the higher-dimensional relationships between traffic flow and speed. Furthermore, through node embedding techniques, the framework autonomously updates a dynamic traffic network graph during model training, thereby simulating real-world traffic conditions. Finally, by integrating graph convolution (GCN) operations in place of the linear transformations within a gated recurrent unit (GRU), the model synchronously captures both temporal and spatial dimensions, effectively revealing complex spatio-temporal dependencies. Experimental evaluations on the widely recognized PEMS series datasets demonstrate that our model surpasses the leading baseline models in medium- to long-term predictions, thereby validating the efficacy and superiority of the proposed hybrid deep learning approach.

Container Throughput Forecasting for Ningbo-Zhoushan Port using Combined SARIMA-LSTM Approach

Presenter: Zhihao Zhou (Changzhou Institute of Technology)
Author(s): Zhihao Zhou;Lijun Wang;Yuting Wu;Zhenhua Yang;Siyuan Miao

Under the guidance of national-level strategies such as the Belt and Road Initiative, the Maritime Power Strategy, and the Yangtze River Economic Belt Development Plan, the significance of port construction has become increasingly apparent. In particular, with the extensive promotion of the Zhejiang Free Trade Zone and the establishment of the National Bulk Commodity Storage and Transportation Base in Ningbo-Zhoushan, new prospects have emerged for the development of Ningbo-Zhoushan Port. Accurate prediction of container throughput holds indispensable importance for future planning, economic analysis, and strategic decision-making in ports. However, there are many factors that affect throughput, and their characteristics need to be considered. Therefore, this article constructs a combined model to predict throughput. The paper first analyzes the features of the container throughput volume sequence of Ningbo-Zhoushan Port from 2012 to 2022, and concludes through one-variable regression that the time series has both linear and nonlinear characteristics. Then, systematic clustering analysis is used to determine that per capita GDP and freight volume are the typical influencing factors of the port's container throughput volume. After that, a seasonal difference self-regression moving average (SARIMA) linear prediction model is built using the training set data to predict the time series based on its temporal features. Then, the prediction residual of the SARIMA model is input together with the selected typical influencing factors into the long short-term memory (LSTM) nonlinear prediction model. Finally, the prediction results of the SARIMA model and the LSTM model are added together to construct a SARIMA-LSTM hybrid model that is consistent with the changing features of the Ningbo-Zhoushan Port container throughput volume sequence. To evaluate the prediction results of the model, MAE, RMSE, and MAPE were used as three evaluation indicators to analyze the prediction results of the model. By comparing the evaluation index results of the combined model with those of two single models, it is concluded that the combined model has higher accuracy of prediction results because of its combination of linear and nonlinear characteristics. This combination model aids port management in short-term container throughput prediction, facilitating timely decisions, refined management, infrastructure improvement, and competitiveness enhancement.

Grey Modelling of Dynamical Systems for Time-series Analysis and Forecasting

Presenter: Wei Baolei (Nanjing University of Science And Technology)
Author(s): Baolei Wei

This talk explores the comprehensive mathematical foundations and systemic model architecture of Grey Forecasting Models (GFMs) within the framework of dynamical systems, encompassing both differential and difference equations. First, we explain the Cumulative Sum Operator, a defining feature of GFMs that sets them apart from other forecasting models. This operator is characterized as an integral operator, which enables the modeling of accumulation phenomena through differential equations. Second, we present a taxonomy of GFMs, systematically categorizing existing models into four distinct scenarios. This classification offers a universal framework for understanding GFMs, facilitating the selection of the most appropriate model for a given problem. The taxonomy not only clarifies the relationships between different GFMs but also highlights their adaptability to various forecasting challenges. Third, we demonstrate the practical impact of GFMs by showcasing their successful applications across diverse fields, including energy, traffic, and steel engineering. Throughout, we emphasize the importance of physical intuition and parameter optimization in advancing the theoretical understanding and practical application of GFMs. By leveraging insights from dynamical systems, we provide a holistic understanding of GFMs, bridging the gap between theoretical foundations and practical applications, while highlighting their potential for future research and innovation.

Grey Forcasting Models for Prognostics Health Management

Presenter: Naiming Xie (Nanjing University of Aeronautics and Astronautics)
Author(s): Naiming Xie

This paper systematically explores the modeling mechanisms of grey forecasting and its applicability to remaining useful life (RUL) prediction. We first employ integral matching to reveal the intrinsic mechanism of generation accumulation, which forms the foundation of grey modeling. Furthermore, we use increments and growth rates to provide an in-depth explanation of the model's dynamic evolution process and structural characteristics. Based on these theoretical insights, we classify and summarize the suitability of grey forecasting models for RUL prediction in diverse engineering scenarios. Specifically, we demonstrate the effectiveness of grey models in fatigue crack propagation prediction and RUL estimation for alloy steel components, online knee-point identification for battery degradation in electric vehicles, battery health prognostics, RUL prediction for lithium-ion batteries, and fatigue life prediction of aircraft lap joint structures. Comparative experiments across these domains show that grey models not only offer high forecasting accuracy but also exhibit strong robustness and adaptability. These results highlight the broad applicability and practical value of grey forecasting models in predictive maintenance and life cycle management tasks.

Fourier Time-varying Grey Model for Seasonal Demand Forecasting

Presenter: Lili Ye (Zhongnan University of Economics and Law)
Author(s): Lili Ye;Naiming Xie;Zhongju Shang

Seasonal demand forecasting is critical for effective retail management. However, conventional forecasting methods face difficulties accurately estimating seasonal variations, owing to time-varying demand trends and limited data availability. In this paper, we propose a Fourier time-varying grey model (FTGM) to tackle this issue. The FTGM builds upon grey models, which are effective with limited data, and leverages Fourier functions to approximate time-varying parameters that allow it to represent seasonal variations. A data-driven selection algorithm adaptively determines the appropriate Fourier order of the FTGM without prior knowledge of data characteristics. Using the well-known M5 competition data, we compare our model with state-of-the-art forecasting methods taken from grey models, statistical methods, and architectures of neural network-based methods. The experimental results show that the FTGM outperforms popular seasonal forecasting methods in terms of standard accuracy metrics, providing a competitive alternative for seasonal demand forecasting in retail companies.

Stationary Density Estimation for ARCH Models with Applications in Quantile Regression

Presenter: Manuel Dario Hernandez-Bejarano (Central Bank of Colombia (Banco de la Republica))
Author(s): Manuel Dario Hernandez-Bejarano;Zhibiao Zhao

This document studies a more efficient density estimation for the stationary density of nonparametric autoregressive conditional heteroscedastic models. These models are important tools in analyzing time series, specifically in economic and financial applications where the goal is modeling and understanding the volatility of the statistical data. We demonstrate that in the presence of an unknown nonparametric variance structure, we can establish the root n consistency of the proposed density estimator, improving the widely used nonparametric kernel density estimator whose convergence rate is inferior. A numerical study confirms the results. The density estimator is applied to the S&P 500 Index data. Finally, we showcase a practical implementation of the proposed density estimator in quantile regression. Specifically, we propose to construct a more accurate estimate of the limiting variance of the estimated coefficients in a quantile regression model whose errors follow a nonparametric autoregressive conditional heteroscedastic structure. Simulation studies show that using the new density estimator leads to a more accurate estimation of the asymptotic variance compared to the results obtained by using the classical density estimator. To illustrate the application of this methodology in estimating the asymptotic variance, we apply it to the monthly inflation rate of the United States.

Communication of Inflation Forecast Uncertainty: The Influence of Probabilistic Communication on Public′ s Perception and Economic Decision-making

Presenter: Nora Luo (Imperial College London)
Author(s): Peiqi (Nora) Luo

Forecasting has developed and evolved over recent decades with applications across a wide range of industries. Forecasts provide estimates about the future, to be used as input to decision-making and risk management. Nevertheless, high forecast accuracy does not necessarily guarantee optimal decision outcomes. Hence, forecast value can only be fully realized when both accurately representing the future and appropriately guiding actions. Randomness in nature, coupled with data limitations and model imperfections, inevitably results in uncertainty within forecasts. Forecast uncertainty is informative as it provides a comprehensive overview of forecaster's best efforts. Yet, it may also be misleading and confusing if people fail to understand it or disregard its worth. In practice, forecasters and decision-makers often occupy different roles, inducing potential communication gaps. An effective communication should enable the audience to accurately understand the information conveyed by forecast and to support rational and analytical reasoning and decision-making. Traditionally, research has been split between enhancing forecast quality and investigating decision-making under uncertainty. Although meteorology has seen progress in studying behavioural responses to probabilistic forecasts, research on such responses in economic and financial contexts remains comparatively limited. As a key macroeconomic indicator, inflation forecasts are communicated to the public and guide economic decisions for various stakeholders. Using inflation forecasts (for single and multiple lead times), this research conducted a survey-based empirical study to evaluate users' ability to interpret embedded uncertainty presented in various formats, including the iconic fan chart. Additionally, it explored, from a behavioural economics perspective, how probabilistic information is internalised and used to shape decisions. The participants are sampled from the UK public and randomised in treatment groups. The interpretation accuracy, reported trust, perceived risk and decision quality were measured and analysed, proving insights into the effectiveness of different forecast communication strategies. The outcome of this research aims to guide policymakers and communicators in refining forecast communication with enhanced clarity and credibility, aligning public expectations of future inflation and empowering public to make informed financial decisions without the risk of misinterpretation.

Time-varying Probabilistic Forecast Combinations Based on Particle Filtering: Diversity Matters.

Presenter: Xiaorui Luo (Beihang University)
Author(s): Xiaorui Luo;Yanfei Kang;Xue Luo

Mitigating uncertainty in probabilistic forecasts requires adaptive strategies to address model redundancy and the dynamic nature of economic conditions. This study introduces a novel framework that integrates diversity regularization into a Bayesian ensemble forecasting system. In this system, time-varying model weights are optimized dynamically using particle filtering. By incorporating real-time diversity metrics, our approach enhances prior knowledge and enables continuous updates of predictive densities while penalizing correlated predictors. This method expands on traditional combination techniques by jointly estimating time-varying biases, inter-model dependencies, and calibration uncertainties within a unified nonlinear filtering framework. Empirical results in macroeconomic and commodity price forecasting demonstrate significant improvements in both point and density accuracy, outperforming individual models and other time-varying combination methods. Additionally, the framework includes diagnostic tools for identifying model incompleteness, further advancing ensemble forecasting in the face of uncertainty.

Interval Forecasting of Tourism Demand Using a Novel Intelligent Combination Technique

Presenter: Guo Honggang (Dongbei University of Finance and Economics)
Author(s): Honggang Guo;Xiaoyu Chen;Doris Chenguang Wu;Yu Jin

Timely demand forecasting plays a vital role in supporting the decision-making of practitioners in the tourism industry. However, existing studies have primarily concentrated on deterministic point forecasts, which fail to account for uncertainty for informed decision-making. On the other side, Combination forecasting has proven to be an effective approach for improving the accuracy of demand forecasts. Nevertheless, a key challenge remains in selecting the optimal subset of individual models from the available models and determining the optimal combination weights. This study therefore innovatively propose an intelligent combination strategy for tourism demand interval prediction, leveraging an improved multi-objective slime mold optimization algorithm. This strategy strengthens the interaction and coupling relationships between individual forecast intervals while adaptively selecting the optimal subset of individual models, eliminating the need to exhaustively test all possible combinations. Furthermore, the multi-objective optimization module within the strategy balances interval coverage and interval width, effectively addressing the inherent limitations of traditional interval forecasting approaches. To empirically validate this novel intelligent combination forecasting strategy, daily tourism demand from three attractions in China are analyzed. The results reveal that (1) the proposed intelligent combination strategy generates more accurate and reliable forecast intervals compared to eight benchmark models, including ETS, ARIMAX, MLP, ELM, RF, XGboost, LSTM and Transformer. (2) The proposal model outperforms both the simple average and two traditional intelligent combination models, including the particle swarm optimization algorithm-based and the gray wolf optimization algorithm-based intelligent combination models. These findings offer valuable insights and novel approaches for improving interval forecasting and supporting informed decision-making in practice. Keywords: Tourism interval forecasting; Combination forecasts; intelligence optimization algorithm; Tourism attraction

A Novel Pattern Recognition Method for Tourism Demand Forecasting

Presenter: Junfeng Huang (Sun Yat-sen University)
Author(s): Junfeng Huang;Doris Chenguang Wu;Mingming Hu

Traditional pattern recognition forecasting methods typically involve searching for and matching slices of time series that resemble the target window within a single series, using subsequent values to predict the target window's future values. These methods have been successfully applied in tourism demand forecasting, yielding promising results. However, challenges arise due to the large data granularity in many tourism datasets and the shifts in tourism demand patterns observed during the COVID-19 pandemic. Under such circumstances, traditional pattern recognition methods often struggle with a lack of sufficient valid data for accurate identification. To address this issue, we innovatively construct a demand time-series pool for pattern recognition by including tourism demand data from multiple regions, aiming to enhance the pattern recognition process. However, the discrepancies in the magnitude of values between the individual series complicate similarity recognition across time series. Moreover, conventional techniques such as normalization tend to alter the inherent patterns (e.g., volatility levels) of the original data significantly. To overcome these challenges, this study proposes a novel forecasting method in which the demand series are transformed into growth rates and time series similarity assessment indices are constructed based on wavelet transform and Manhattan distance. This approach also considers both trend and fluctuation similarities, enabling the

identification of similarity windows of varying lengths across the time series pool. Daily inbound tourism demand from multiple countries and regions to Macau are used for empirical examination. The results demonstrate that this proposed novel pattern recognition method outperforms commonly used baseline models (e.g., SARIMA, NAÏVE, and ETS) and the traditional single-series pattern recognition approach. The findings of this study enrich pattern recognition forecasting literature by methodological innovation.

Development of a Real-time Tourism Demand Forecasting System for Hong Kong

Presenter: Xinyan Zhang (The Hong Kong Polytechnic University)
Author(s): Xinyan Zhang;Haiyan Song;Hanyuan Zhang

Abstract: Real-time forecasting, also known as nowcasting, involves predicting economic activities in the current time horizon. Interest in tourism demand nowcasting has increased in recent years due to the growing uncertainty in tourism activity. Nowcasting incorporates real-time data into short-term forecasts, providing valuable data-driven information to researchers, managers, and policymakers in the tourism sector. This study innovatively developed a cloud-based real-time tourism demand forecasting system to nowcast mainland Chinese visitor arrivals to Hong Kong during the tourism recovery stage. Baseline nowcasts were generated by combining survey-based business and consumer confidence index data with real-time search query volume data in a model. Given the high uncertainties in tourism recovery, the model-generated nowcasts for the current quarters were further adjusted by a panel of experts through the Delphi method within the system. The system also provided results from quarterly tourist confidence surveys conducted in mainland Greater Bay Area (GBA) cities to support the experts during the Delphi process. This adjustment was made to incorporate important extra information into the final nowcasts and improve forecasting accuracy. The results provided strong evidence of the outstanding ex-ante nowcasting performance of the model, and the significant value added by human judgment to the nowcast accuracy. Keywords: tourism demand nowcasting, forecasting support system, cloud-based system, judgmental forecasting, Delphi method Acknowledgements: This study is fully supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. UGC/FDS24/B02/22).

Forecasting the El Niño Southern Oscillation (ENSO) Index with High-Dimensional Factor Models

Presenter: Alessandro Giovannelli (University of L'Aquila)
Author(s): Alessandro Giovannelli;Tommaso Proietti

The El Nino/Southern Oscillation (ENSO) is a major driver of interannual climate variability. This study introduces a novel methodology, based on large dynamic factor models (DFMs), aimed at improving the prediction of sea surface temperatures (SST) in the four El Nino regions, key indicators of ENSO activity. A key issue in DFMs lies in selecting relevant variables for factor extraction, as the inclusion of noisy or uninformative series can degrade predictive performance. While existing approaches use static correlation to identify targeted principal components, this overlooks variables that may have weak static but strong frequency-specific relationships with the target. We propose a coherence-based variable selection method that identifies predictors with significant comovement with the target variable within the most relevant frequency bands. This allows for a more effective reduction of noise and improved factor estimation. Applying our method to a high-dimensional SST dataset, with El Nino 3.4 as the target, we perform a real-time forecasting exercise and compare results with models using complete dataset. Preliminary results indicate that coherence-based selection offers potential to enhance the accuracy of ENSO predictions.

Integrating Renewable Energy Forecasting and Sustainable Development Goals: Proposal of a Framework for Plant Site Selection

Presenter: Gustavo Melo (Pontifical Catholic University of Rio de Janeiro (PUC-Rio))
Author(s): Gustavo Melo;Fernando Luiz Cyrino Oliveira;Paula Maçaira;Bernardo Bessada

The transition to renewable energy is a key pillar for achieving several Sustainable Development Goals (SDG) and mitigating global socioeconomic and environmental crises. In this context, selecting sites for renewable power

plants is a complex decision-making process that must consider multiple criteria, such as energy production potential and socio-environmental factors. In this process, modeling the stochastic nature of renewable energy sources (RES) is essential to capture the variability of future energy supply. Additionally, considering criteria directly related to the SDG targets is crucial to fulfilling the commitments of the 2030 Agenda of the Paris Agreement. In this sense, this study proposes a framework to generate a probabilistic indicator to support the site selection process for renewable energy farms based on RES forecasts and SDG-related criteria. An innovative forecasting method is proposed, correlating the predicted renewable source of different sites. The developed method follows a hybrid approach, combining machine learning techniques, Markov chains, and Monte Carlo simulation. Then, the forecasts are combined with SDG variables through multi-criteria decision-making techniques (MCDM), obtaining a probabilistic indicator to rank a set of sites. A case study was conducted in the Brazilian state of Rio Grande do Norte for the framework's validation, focusing on ranking municipalities for wind plants. Regarding the results, the forecasts captured the main properties of the RES of all the sites, including out-of-sample periods. Furthermore, the probabilistic indicator demonstrated that several municipalities without existing wind farms but with high production potential could benefit from future installations. Therefore, using the indicator in site selection could foster regional development and contribute to sustainability goals. It is important to highlight that the proposed framework can be applied to other regions, renewable energy sources, and criteria, provided data is available.

Cyclical Time Series: an Empirical Analysis of Temperatures in Central England over Three Centuries

Presenter: Liudas Giraitis (Queen Mary University of London)
Author(s): Liudas Giraitis;Fulvia Marotta;Peter CB Phillips

This paper builds on methodology that corrects for irregular spacing between realizations of unevenly spaced time series and provides appropriately corrected estimates of autoregressive model parameters. Using these methods for dealing with missing data, we develop time series tools for forecasting and estimation of autoregressions with cyclically varying parameters in which periodicity is assumed. To illustrate the robustness and flexibility of the methodology, an application is conducted to model daily temperature data. The approach helps to uncover cyclical (daily as well as annual) patterns in the data without imposing restrictive assumptions. Using the Central England Temperature (CET) time series (1772 - present) we find with a high level of accuracy that temperature intra-year averages and persistence have increased in the later sample 1850-2020 compared to 1772 - 1850, especially for the winter months, whereas the estimated variance of the random shocks in the autoregression seems to have decreased over time.

Low-rank Tensor Autoregressive Models for Mortality Modeling

Presenter: Tim Boonen (The University of Hong Kong)
Author(s): Tim Boonen;Yuhuai Chen

Mortality patterns in closely related sub-populations often exhibit similarities, suggesting that mortality forecasts for individual sub-populations could be enhanced by incorporating patterns from larger related groups. In this paper, we focus on multiway multi-population mortality modeling, in which there exists a multiway mortality array comprising mortality rates of populations disaggregated by various socio-demographic attributes, such as gender, age, smoking/non-smoking, and country or region. Each dimension of the array corresponds to one attribute. We propose a tensor autoregressive (TAR) model to efficiently model and forecast such multiway mortality arrays. Unlike existing vector autoregressive (VAR) models, the TAR model preserves the multiway structure, more effectively incorporating patterns across groups and attributes. By utilizing the CANDECOMP/PARAFAC (CP) and Tucker decomposition, the proposed TAR model captures underlying low-dimensional tensor dynamics, enabling significant dimension reduction and clustering of groups sharing similar mortality dynamics. Finally, an empirical analysis using three-way mortality data (Age, Country, and Gender) demonstrates that the model achieves strong in-sample fit and satisfactory out-of-sample forecasting performance. Furthermore, we demonstrate that coherence can be easily imposed under the CP decomposition.

A Functional Signature Approach for Mortality Forecasting

Presenter: Zhong Jing Yap (University of Malaya)
Author(s): Zhong Jing Yap;Dharini Pathmanathan;Sophie Dabo-Niang

Signatures from rough path theory are used to summarise paths with a compact representation of their respective iterated integrals. It naturally captures linear functionals essential for predicting its effects on systems, which makes it an appealing feature map for statistical learning and regression tasks. In this work, an integration of truncated and randomized signature methods into the framework of the Hyndman-Ullah (HU) model is introduced to model and forecast mortality rates. We first demonstrate the proposed signature-embedded models on whole populations, followed by subpopulations by employing grouped functional times series methods. Preliminary findings show that the randomized signature framework performs better than its truncated counterpart on Japanese mortality data.

Constructing Prediction Intervals for the Age Distribution of Deaths

Presenter: Hanlin Shang (Macquarie University)
Author(s): Hanlin Shang;Steven Haberman

We introduce a general procedure to construct prediction intervals for the age distribution of deaths. The age distribution of deaths is an example of constrained data, which are nonnegative and have a constrained integral. A centered log-ratio transformation and a cumulative distribution function transformation are used to remove the two constraints, where the latter transformation can also handle the presence of zero counts. Our general procedure divides data samples into training, validation, and testing sets. Within the validation set, we can select an optimal tuning parameter by calibrating the empirical coverage probabilities to be close to their nominal ones. With the selected optimal tuning parameter, we then construct the pointwise prediction intervals using the same models for the holdout data in the testing set. Using Japanese age- and sex-specific life-table death counts, we assess and evaluate the interval forecast accuracy with a suite of functional time-series models.

Research on a Multi-Agent Collaborative Early Warning System for Multinational Corporate Interest Security Based on Reinforcement Learning

Presenter: Xiaowei Li (China University of Petroleum (East China))
Author(s): Xiaowei LI;Xiaofeng XU

Against the backdrop of deepening globalization, safeguarding multinational corporate interests faces three major challenges: escalating geopolitical tensions, intensified fluctuations in the industrial supply chains, and frequent social and cultural conflicts. This research proposes a Multi-Agent Security Early Warning System (MAS-ESWS) that integrates Multi-Agent Reinforcement Learning (MARL) and spatiotemporal cognitive computing. It establishes an intelligent architecture of "data-model-decision" co-evolution, and addresses the two major technical bottlenecks of dynamic strategy coupling instability and cross-modal semantic gap. MAS-ESWS follows the "Centralized Training-Decentralized Execution" (CTDE) paradigm and integrates four specialized agents: the intelligence collection agent optimizes the multi-source data access strategies to improve the accuracy and timeliness of intelligence acquisition; the risk assessment agent combines time series analysis and graph convolutional networks to construct a dynamic coupling model of multi-dimensional influencing factors to achieve accurate early warning; the decision arbitration agent optimizes the warning level decision based on differentiable game theory to reduce the misjudgment rate; the emergency response agent uses the reinforcement learning framework to optimize resource scheduling and improve response efficiency. This research introduces a multi-agent collaborative game equilibrium mechanism that combines local observation constraints and global reward reshaping to mitigate strategy oscillations in dynamic environments. It develops a spatiotemporal attention-driven multimodal fusion framework to bridge the semantic gap in heterogeneous data. Furthermore, the proposed model undergoes collaborative validation in cross-border risk prevention scenarios, ensuring both theoretical rigor and practical applicability, providing theoretical support and practical paths for risk prevention, control, and decision-making optimization. Empirical results confirm the system's superiority in risk identification accuracy, environmental adaptability, and decision robustness, establishing a novel paradigm for safeguarding multinational corporate interests in complex globalized contexts.

Transformer-based Downside Risk Forecasting: A Data-Driven Approach with Realized Downward Semi-Variance

Presenter: Yuetong Zhang (Shandong University)
Author(s): Yuetong Zhang

Realized downward semi-variance (RDS) has been realized as a key indicator to measure the downside risk of asset prices, and the accurate prediction of RDS can effectively guide traders' investment behavior and avoid the impact of market fluctuations caused by price declines. In this paper, the RDS rolling prediction performance of the traditional econometric model, machine learning model, and deep learning model is discussed in combination with various relevant influencing factors, and the sensitivity analysis is further carried out with the rolling window length, prediction length, and a variety of evaluation methods. The empirical results show that: 1) When the rolling window is set to 20, the overall prediction effect of the model in this paper is the best. Taking the Transformer model as an example, compared with the prediction results under the rolling window length of 5, 10, and 30, the RMSE improvement ratio reaches 24.69%, 15.90%, and 43.60%, respectively. 2) The multi-variable Transformer model shows a better forecasting effect. Compared with traditional econometric, machine learning, and deep learning models, the average increase percentage of RMSE, MAE, MAPE, and SMAPE indicators is 52.23%, 20.03%, 62.33%, and 60.33%, respectively. 3) In multi-step prediction scenarios, the DM test statistic of the Transformer model is significantly positive, and the prediction accuracy of the Transformer model remains stable as the number of prediction steps increases. An in-depth study of RDS forecasting is of great value to capture the characteristics of downside risks, enrich the financial risk measurement system, and better evaluate potential losses.

Forecasting Probability Distributions of Financial Returns with Deep Neural Networks

Presenter: Jakub Michańków (University of Warsaw)
Author(s): Jakub Michańków

This study examines the application of deep neural networks for forecasting the probability distributions of financial returns. Instead of focusing only on the point predictions, the proposed approach uses deep learning architectures to forecast the full predictive distribution of asset returns, by estimating the parameters of three types of probability distributions: normal, Student's t, and skewed Student's t. The research examines models based on recurrent (LSTM) and convolutional (CNN) networks, designed to capture non-linear dependencies and complex patterns in financial time series. The performance of these models is evaluated using probabilistic forecast accuracy metrics such as Log Predictive Score (LPS), Continuous Ranked Probability Score (CRPS) and Probability Integral Transform (PIT), as well as their implications for financial risk estimation, including Value-at-Risk and Expected Shortfall. The results indicate the potential and limitations of deep learning in probabilistic forecasting, offering perspectives on its practical applications in risk management and investment strategies.

The Puzzle of Overpriced European Treasury Auctions

Presenter: Jose Faias (Universidade Catolica Portuguesa)
Author(s): Jose Faias;Jose Miguel Cardoso-Costa;Patrick Herb;Mark Wu

Underpricing in Treasury auctions has been extensively examined, but recent evidence suggests that the primary dealership model used in many euro area countries leads to overpriced auctions. Analyzing 15 years of Portuguese Treasury auctions using a unique dataset that allows tracking dealers over time, we show that more aggressive bidding behavior is directly linked to dealers seeking access to syndication fees and post-auction non-competitive offerings, resulting in consistently overpriced auctions. Furthermore, we have found that this institutional bundling design allows a specific group of dealers to effectively turn the anticipated losses in the competitive auction market into post-auction profits.

Conditionally Optimal HAR Forecasts

Presenter: Andrey Vasnev (University of Sydney Business School)
Author(s): Andrey Vasnev;Adam Clements;James Tannock

This study advances the investigation of the heterogeneous autoregressive (HAR) model through the lens of forecast combination, following the approach of Clements and Vasnev (2024). In order to enhance the prediction of realized volatility, we apply the conditionally optimal framework recently developed by Gibbs and Vasnev (2024). Traditional HAR models and their univariate components often fail to adapt swiftly to large rapid market shifts, such as those seen during the Global Financial Crisis (GFC) in 2008 and the more recent COVID-19 pandemic. By explicitly modeling the conditional bias of the univariate model forecasts with simple autoregressive structures, we extract valuable information that can be included in the conditionally optimal weights. This method yields substantial gains in predictive accuracy across the S&P 500 and Dow Jones indices, with the most pronounced benefits observed at longer forecast horizons.

Forecasting Expected Stock Returns: a Shrinkage Combination Approach

Presenter: Han Feng (University of Chinese Academy of Sciences)
Author(s): Han Feng;Difang Huang;Yinuo Ren;Hanfeng Zhou;Pengyang Song

Forecast combination has gained popularity in recent years. Popular combination schemes have evolved from equal weights (EWs), which utilize simple averages, to optimal weights (OWs), which optimize weights by minimizing mean squared error (MSE). Extensive research has not only validated the robustness of EWs but also introduced the concept of shrinkage, shrinking OWs towards EWs. In this paper, we study the predictability of a shrinkage combination method for stock returns, using 13 financial and macroeconomic predictors in total. This combination model exhibits significant forecasting power both in and out of sample, with out-of-sample annual R^2 2.30%, greater than individual predictors and other benchmarks. The shrinkage combination also generates significant economic gains for a mean-variance investor with a utility gain of over 400 basis points and a Sharpe ratio of 0.91 per annum. An investigation of elements in shrinkage weights shows that the source of predictability may come from the DP predictor.

Human Expertise vs. Machine Learning: Forecasting China′ s Outbound Tourism Through a Two-Round Delphi Survey

Presenter: Shanshan (Vera) Lin (School of Management, Zhejiang University, China)
Author(s): Shanshan (Vera) Lin;Hanyuan Zhang;Haiyan Song

This study evaluates the comparative efficacy of the expert-based Delphi method and quantitative forecasting models in predicting Chinese outbound tourism demand (2023Q1–2027Q4) under conditions of high market uncertainty. Employing a two-round Delphi survey, forecasts were generated for five key destinations: Hong Kong, Macau, Japan, Singapore, and Thailand. A panel of 129 industry professionals from hotels, OTAs, MICE, airlines, theme parks, government tourism authorities, and consulting sectors participated in the first round (March 29–June 23, 2023) survey, providing baseline predictions via Excel-based questionnaires. In the second round (December 20, 2024–February 19, 2025), 102 respondents revised their forecasts after reviewing aggregated forecast results, individual forecast errors (calculated using MAPE, MSE, and MAE metrics), and actual tourism data covering 2023Q1–2024Q3. Participants were further encouraged to refine their predictions using large language models (LLMs). The findings offer critical insights into the relative performance of expert-driven versus quantitative forecasting approaches in volatile conditions, while also projecting detailed trends of Chinese outbound tourism. This research provides actionable guidance for industry stakeholders seeking to enhance strategic decision-making amid uncertainty.

Collaborative Multi-agent Strategies and Tourism Demand Forecasting

Presenter: Peiying Zhang (Sun Yat-sen University)
Author(s): Peiying Zhang;Doris Chenguang Wu

Abstract: LLM-based agents have demonstrated exceptional performance across diverse tasks such as text generation and open-domain question answering. However, their potential in forecasting tourism demand remains an

underexplored frontier. This study pioneers the systematic examination of LLM-based agents' tourism demand prediction capabilities under different collaborative strategies. This exanimation contains three folds. First, we construct three different AI agents based on different LLMs to compare the prediction capabilities of a single agent operating independently versus those collaborating through parallel cooperation. Second, by incorporating differentiated identity profiles into the same three LLMs, we establish three agent groups, each of which contains agents with three identities: data analyst, tourism market researcher, and tourism economist, and then compare the intra-group parallel collaboration with the inter-group parallel collaboration in the previous phase. Thirdly, we implement a chain communication strategy as an alternative to parallel cooperation, conducting comparative evaluations of both collaborative paradigms. Hong Kong daily inbound tourism arrivals dataset are adopted for empirical validation. The results demonstrate that parallel collaboration outperforms both independent agent operations and chain-based cooperation. Our findings not only validate the exceptional capabilities of LLM-based agents in tourism demand forecasting but also identify optimal collaborative strategies for multi-agents, offering implementable insights for tourism authorities to enhance crowd management through AI-powered predictive analytics. Keywords: tourism demand forecasting; LLM-based agents; parallel cooperation; chain communication strategy; artificial intelligence

Forecasting Tourist Arrival Using Image Data and LLM

Presenter: Haoran Zhang (Hunan Normal University)
Author(s): Kaijian He; Linyuan Zheng; Ying Lu; Yingchao Zou; Haoran Zhang

In this paper we have proposed a new tourist arrival forecasting dmodel based on image data and the large language model. It takes advantage of both the larger data set based on image data and the larger model. Big dataset has been constructed beyond the traditional time series data to image data and textual data such as review comments and newsa. We introduce the larger language model to directly aligned and combines these diverse data in different modals. It takes advantage of the rich information in these diverse multi modal data using the powerful attention mechanism and the cross modal embedding in large language models. A series of different large language models such as deepseek can be used to model the multimodal data. Empirical studies have been conducted using macau as the test ground for the model evaluation. Experiment results confirm the superior forecasting accuracy from the introduction of a bigger data set and large language model. Comprehensive experiment evaluation show that inclusion of textual and image data set leads to the improved forecasting accuracy. The size of the large language model also have critical impacts on the forecasting accuracy as the larger size model have more parameters to model exponentially increasing variable connections and heterneous information in the data set.

Virtue of Complexity in Housing Price Prediction

Presenter: Jinfeng Gan (Fudan University)
Author(s): Jinfeng Gan;Zaichao Du;Donghao Zhang

This study employs panel data from ten Canadian provinces spanning from January 2005 to August 2024 to introduce the "Virtue of Complexity" theory into the field of housing price prediction, elucidating the applicability boundaries of complexity-driven models in low-frequency trading markets. Addressing the limitations of traditional hedonic price models in capturing spatiotemporal nonlinear characteristics, we construct a province-specific time-series regression framework that incorporates spatiotemporal dependencies. This framework is combined with Random Features and Nyström kernel approximations to achieve high-dimensional nonlinear representations. Furthermore, the study systematically evaluates the dimensionality gain achieved through Gaussian Noise Injection techniques and validates the universality of the complexity-driven model using Random Forest models. The findings reveal that: (1) Complexity-driven models based on kernel approximations exhibit a double descent phenomenon in the over-parameterized region, but their low-frequency signal characteristics result in a significantly weaker improvement in predictive performance compared to financial markets; (2) The choice of kernel function type has a substantial impact on predictive stability. The ReLU kernel demonstrates superior stability in capturing policy shocks and long-term trends, whereas the RBF and Sigmoid kernels exhibit greater volatility; (3) Gaussian Noise Injection is locally effective in scenarios without feature selection, but its combination with feature selection tends to introduce interfering signals; and (4) Random Forest models generally exhibit a single descent phenomenon in non-stationary data, but a few provinces show signs of overfitting, highlighting the significant constraints imposed

by data distribution heterogeneity on the selection threshold for model complexity. Our results refine the universality of the "Virtue of Complexity" theory in low-frequency economic contexts and provides valuable guidance for model optimization in housing market prediction.

Somewhere between Simplicity and Complexity

Presenter: Yin Zihan (Fudan University)
Author(s): Zihan Yin;Zaichao Du;O-Chia Chuang

Recently, the phenomena of "double-descent" and "virtue of complexity" have gained popularity in statistics and economics, where people find that models with number of parameters far exceeding number of observations perform well in forecasting. In this paper, we also document better return forecasts with more parameters using some common machine learners in different setups. However, we generally find that a Nyström method with a number of parameters 1/6 of the number of observations, as well as a ridge regression with a large penalization parameter and a handful of regressors, outperforms the complex models. Our results cast some doubt on the necessity of using complex models and instead recommend the use of the Nyström method in return prediction.

How Predictable is Chinese Stock Market? New Comprehensive Empirical Evidence

Presenter: Ting Zhang (University of International Business and Economics)
Author(s): Ting Zhang;Haibin Xie

A comprehensive empirical study is performed on the Chinese stock market to reevaluate its predictability. We find the predictability of the Chinese stock market is highly related to a bubble period, from 2006 to 2008. Once the bubble period is removed, neither macroeconomic variables nor technical indicators are found to be informative. This finding is in line with Goyal and Welch (2008, 2024), who find structural breaks in the economy can critically undermine the predictability of financial markets. This paper provides new empirical evidence confirming the difficulty of predicting financial markets.

Decoupling Systemic Risk into Endopathic and Exopathic Competing Risks Through Autoregressive Conditional Accelerated Fréchet Model

Presenter: Ji Jingyu (Capital University of Economics and Business)
Author(s): Jingyu Ji;Deyuan Li;Zhengjun Zhang

Identifying systemic risk patterns in geopolitical, economic, financial, environmental, transportation, epidemiological systems and their impacts is the key to risk management. This paper proposes a new nonlinear time series model: autoregressive conditional accelerated Frechet (AcAF) model and introduces two new endopathic and exopathic competing risk indices for better learning risk patterns, decoupling systemic risk, and making better risk management. The paper establishes the probabilistic properties of stationarity and ergodicity of the AcAF model. Statistical inference is developed through conditional maximum likelihood estimation. The consistency and asymptotic normality of the estimators are derived. Simulation demonstrates the efficiency of the proposed estimators and the AcAF model's flexibility in modeling heterogeneous data. Empirical studies on the stock returns in S&P 500 and the cryptocurrency trading show the superior performance of the proposed model in terms of the identified risk patterns, endopathic and exopathic competing risks, being informative with greater interpretability, enhancing the understanding of the systemic risks of a market and their causes, and making better risk management possible.

Sparse Multivariate Autoregressive Conditional Fréchet Models for High-Frequency Extreme Risk Dynamics

Presenter: Zhengjun Zhang (University of Chinese Academy of Sciences)
Author(s): Zhengjun Zhang;Yongjoon Kim

We introduce a novel class of multivariate Fréchet models that integrate dynamic cross-sectional and temporal

dependencies to effectively capture time-varying marginal characteristics. Building on the classical multivariate maxima of moving maxima (M4) framework, we propose a sparse variant of the M4 random coefficient model (vSM4R) to address estimation challenges in high-dimensional settings. For marginal processes, we extend the autoregressive conditional Fréchet (AcF) model to an mmm-dependent structure (mAcF), enabling richer dependency patterns while maintaining a direct link to the vSM4R through standardization. We establish key probabilistic properties and develop composite likelihood estimation methods with proven statistical guarantees. Through simulations and high-frequency cryptocurrency return data, we demonstrate the robustness of our approach in capturing tail risks and modeling extreme co-movements in financial markets. Our results underscore the necessity of flexible multivariate extreme value models for risk management in volatile environments.

Competing-risk Weibull Survival Model with Multiple Causes

Presenter: Chengxiu Ling (西交利物浦大学慧湖药学院)
Author(s): Chengxiu Ling;Zhengjun Zhang;Kai Wang;Yuqin Mu

The failure of a system can result from the simultaneous effects of multiple causes, where assigning a specific cause may be inappropriate or unavailable. Examples include contributing causes of death in epidemiology and the aetiology of neurodegenerative diseases (e.g., Alzheimer's disease). We propose a parametric Weibull accelerated failure time model for multiple causes, incorporating a data-driven, individualized, and time-varying winning probability (relative importance) matrix. This approach maintains linear structures in competing risks for interpretability, allows for modelling non-Weibull distributed time, and includes penalization for variable selection sparsity. Using maximum likelihood estimation and the expectation-maximization (EM) algorithm, our approach enables simultaneous estimation of regression coefficients and relative cause importance, ensuring consistency and asymptotic normality. A simulation study and an application to survival status from mild cognitive impairment (MCI) to Alzheimer's disease diagnosis demonstrate its effectiveness in addressing cause-mixture problems and identifying informative biomarker combinations, with comparisons to Weibull and Cox proportional hazards models.