Transregio 391

technische universität dortmund

TRR 391
TC02 Workshop

Spatio-temporal Statistics in High Dimensions

November 27 and November 28, 2025

1. Workshop Location
2. Dinner Location
3. Program
4. Abstracts of Contributed Talks
5. Abstracts of Invited Talks

1. Workshop Location

The workshop will take place at TU Dortmund University in the rooms of the Faculty of **Biological and Chemical Engineering**, in the **BCI Gallery**. It is located in the **basement** of the **BCI building**.

Address: TU Dortmund, BCI-Gallery, Emil-Figge Str. 66, 44227 Dortmund

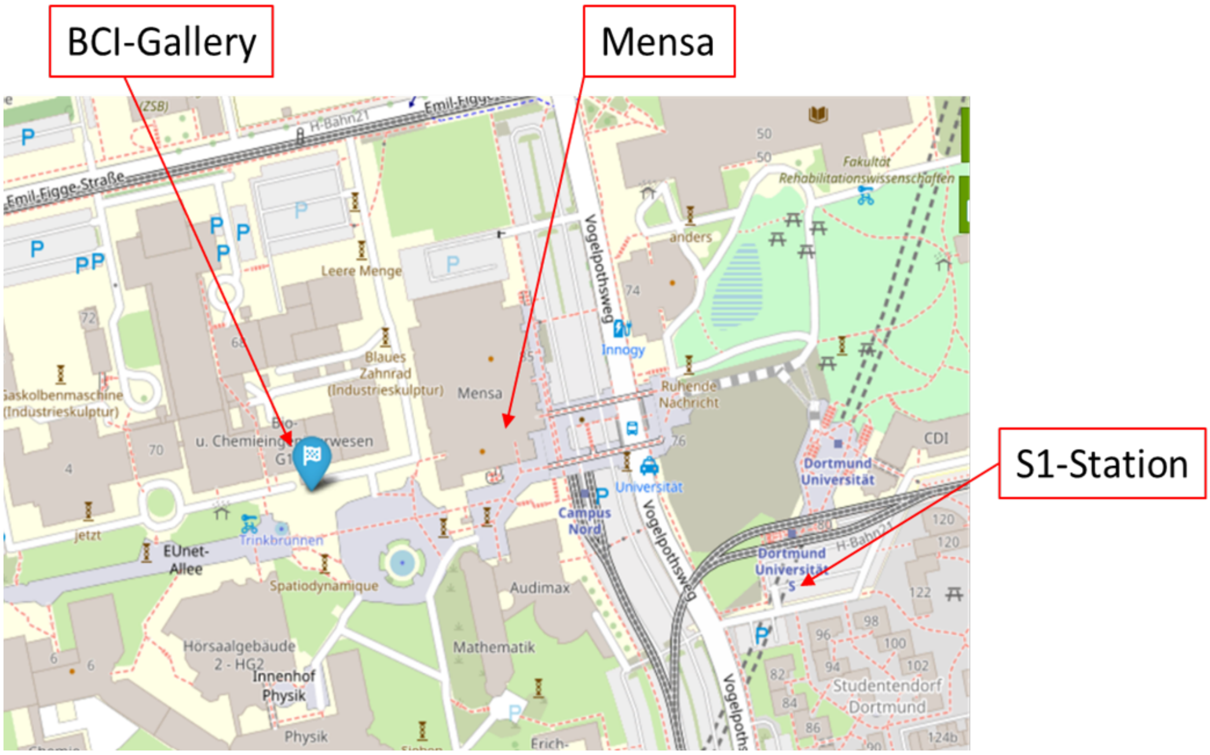

The walking distance from the S1-Station “Dortmund Universität” is about 5-10 minutes. The S1 needs approximately 5 minutes from the railway station ”Dortmund Hbf” and runs every 15 minutes on weekdays.

For further information:
https://www.tu-dortmund.de/en/campus/contact-and-how-to-get-to-the-campus/

2. Dinner Location

We will have a dinner gathering on November 27, at 18:30, for all those who have already registered for the dinner during the registration process.

Please note that this will be a self-payment event.

Address: Maximilian Dortmund, Markt 10, 44137 Dortmund

Transregio 391

technische universität dortmund

3. Program

Thursday	**November 27, 2025**
	Chair: *Florian Ziel*
9:00 - 9:30	Invited: Integrating Complex Covariate Transformations in Generalized Additive Models *Claudia Collarin*
09:35 - 10:10	Invited: Inferring the spatio-temporal structure of climate variability and its external and internal drivers *Nils Weitzel*
10:10 - 10:45	Invited: High-Dimensional Moore-Penrose and Ridge-Type Regularization *Nestor Parolya*
10:45 - 11.00	Coffee and Tea
	Chair: *Axel Bücher*
11:00 - 11:25	Tapered covariance matrix estimation for lattice processes *Antonio Fioravanti*
11:25 - 11:50	Regularized Online Distributional Regression *Florian Ziel*
11:50 - 12:15	Regularization in generalized linear mixed models *Jacob Grytzka*
12:15 - 12:40	Uniform variance reduced simultaneous inference of time-varying correlation networks *Lujia Bai*
12:40 - 14:00	Lunch
	Chair: *Nils Weitzel*
14:00 - 14:35	Invited: D-Vine Copula based Probabilistic Weather Forecasting *Annette Möller*
14:35 - 15:10	Invited: Throwing Vines at the Wall: Structure Learning via Bias-Corrected Random Search *Thibault Vatter*
15:10 - 15:35	Estimating Structure Factor Models with Hidden Factors for Extremes *Alexis Boulin*
15:35 - 16:05	Coffee and Tea
	Chair: *Andreas Groll*
16:05 - 16:30	Efficient score matching based on diffusion models *Benedikt Lütke Schwienhorst*
16:30 - 16:55	Towards Convex Stochastic AC OPF for Radial Distribution Grids *Oleksii Molodchyk*
16:55 - 17:20	The Field Equations of Penalized non-Parametric Regression *Sven Pappert*
from 18:30	Dinner at Maximilian *Markt 10, 44137 Dortmund*

Friday	November 28, 2025
	Chair: *Christoph Hanck*
09:00 - 09:35	Invited: Sparse High-Dimensional Vector Autoregressive Bootstrap *Stephan Smeekes*
09:35 - 10:10	Invited: Spatio-Temporal Count Autoregression *Kostas Fokianos*
10:10 - 10:45	Invited: Ultra-High Dimensional Cointegration *Melanie Schienle*
10:45 - 11:00	Coffee and Tea
	Chair: *Carsten Jentsch*
11:00 - 11:25	The Vector Conditional Autoregressive Wishart Model for Multivariate Stock Market Volatility *Jan Vogler*
11:25 - 11:50	Learning of Smoothing Parameters in GAM Frameworks *Monika Zimmermann*
11:50 - 12:15	Pivotal inference for linear predictions in stationary processes *Sebastian Kühnert*

4. Abstracts of Contributed Talks

Uniform variance reduced simultaneous inference of time-varying correlation networks

Lujia Bai

This paper proposes a unified framework for inferring large-scale time-varying correlation networks via data-driven time-varying thresholds that can control uncertainty simultaneously. The framework allows the dimension of time series vectors to be fixed or diverging at a high polynomial rate of the sample size. It also allows the time series to exhibit changing temporal characteristics beyond stationarity without specific structural assumptions. Motivated by the practical issue that the confidence band of non-parametric estimators of correlations can exceed their natural domain [-1, 1] (see, for example [1])), we propose a simple uniform variance reduction technique. When applied to the construction of a correlation network, the new device yields more accurate thresholds, which enhance the probability of recovering the time-varying network structures. We broaden the applicability of our method by developing difference-based estimators of cross-correlations that are robust to structure breaks in the time-varying mean functions, and by allowing both a fixed and a diverging number of lags in the correlation functions. We prove the asymptotic validity of the proposed method, especially in achieving accurate family-wise error control when disclosing flexible time-varying network structures. The effectiveness of our method in finite samples is demonstrated through simulation studies and data analysis.

Estimating Structure Factor Models with Hidden Factors for Extremes

Alexis Boulin

This work proposes a novel method for estimating the linear factor model $X = AZ$, designed for settings where the latent factors Z are possibly dependent and exhibit extremal behavior. Our approach identifies the loading matrix A and the number of factors K using the tail pairwise dependence matrix—a measure analogous to the covariance matrix but tailored for extremes. This method requires the existence of pure variables, meaning components of X that load exclusively on a single factor. Our proofs for model identifiability are constructive and lead to two practical algorithms for estimating K and A. We demonstrate the method's utility on several real-world datasets, including river discharge data monitored at multiple stations along the Danube, which we use to estimate the probabilities of concurrent extreme events across numerous locations. We demonstrate this method on several real-world datasets, including river discharge data from multiple stations along the Danube, and use it to estimate the probabilities of concurrent extreme events across numerous locations.

Tapered covariance matrix estimation for lattice processes

Antonio Fioravanti

For a stationary lattice process on a regular grid, we consider the estimation of the whole covariance function of the data and generalize the tapering estimation proposed by McMurry and Politis (2010) and Jentsch and Politis (2015) developed for univariate and multivariate time series, respectively, to general lattice process. We prove the consistency of the estimation with respect to the spectral norm and discuss computational challanges that arise from high dimensionality. To achieve efficiency gains, we exploit separability of the covariance function and examine simulations with different levels of spatial dependency. Furthermore, we propose an alternative tapered estimator for separable covariance functions.

Regularization in generalized linear mixed models

Jacob Grytzka

LASSO penalization is a well-established method for variable selection and regularization in regression modeling. However, its application in mixed models, which combine fixed and random effects, presents specific challenges related to the estimation methods used for these models. In our work, we focus on penalizing the fixed effects while leaving the random effects unaffected to preserve their hierarchical variance structures. This approach aims at eliminating irrelevant predictors and reducing model complexity without restricting the flexibility in modeling hierarchical structure. Instead of the classical LASSO penalty, we use a quadratic approximation, ensuring differentiability. This modification facilitates optimization while retaining the core idea of sparse modeling. We implement our LASSO penalty in two widely applied R packages for mixed models. Using extensive simulations, we demonstrate that the new penalty enables efficient variable selection in mixed models even in high-dimensional settings with much more predictors than observations.

Pivotal inference for linear predictions in stationary processes

Sebastian Kühnert

In this paper we develop pivotal inference for the final (FPE) and relative final prediction error (RFPE) of linear forecasts in stationary processes. Our approach is based on a novel self-normalizing technique and avoids the estimation of the asymptotic variances of the empirical autocovariances. We provide pivotal confidence intervals for the (R)FPE, develop estimates for the minimal order of a linear prediction that is required to obtain a prespecified forecasting accuracy and also propose (pivotal) statistical tests for the hypotheses that the (R)FPE exceeds a given threshold. Additionally, we provide new (pivotal) inference tools for the partial autocorrelation, which do not require the assumption of an autoregressive process.

Efficient score matching based on diffusion models

Benedikt Lütke Schwienhorst

We investigate the statistical properties of a score matching variant based on diffusion models. While vanilla score matching can be statistically inefficient for the estimation of multimodal distributions, diffusion models can efficiently sample from multimodal distributions, such as images. We aim to show that the diffusion model-based score matching estimator overcomes the difficulties of vanilla score matching and efficiently estimates multimodal distributions.

Towards Convex Stochastic AC OPF for Radial Distribution Grids

Oleksii Molodchyk

Consideration of stochastic uncertainties is important in power grid operation as it helps to account for the intermittent nature of renewables and/or complicated behaviors of power consumers. One of the methods to solve stochastic AC optimal power flow (OPF) problems available in recent literature is based on polynomial chaos expansions (PCEs). Although this approach provides high-accuracy interpretable results, it can become computationally expensive, especially when applied to large grids with many different uncertainty sources. In this paper, we show that for radial grids, the performance of the PCE-based method can be improved. By exploiting radial structures and using a specifically tailored polynomial basis, we formulate a convex relaxation for the PCE-overloaded stochastic AC OPF. The result is a second-order cone program that can be solved efficiently. We demonstrate the efficacy of our approach on a 47-bus example.

The Field Equations of Penalized non-Parametric Regression

Sven Pappert

In high-dimensional regression settings, penalising a risk is a popular approach for regularisation. In this talk, we examine penalised risks from the perspective of the calculus of variations. We consider risks comprising a fitness term (e.g. MSE) and a gradient-based penalty. We establish the Euler–Lagrange field equations as a systematic approach to finding minimisers of risks involving only first derivatives. We proceed to demonstrate this approach with an example of the MSE penalised by the integral of the squared l2-norm of the gradient of the regression function. The minimiser of this risk is the solution to a second-order inhomogeneous PDE, where the inhomogeneity is the conditional expectation of the target variable given the features. We discuss the properties of the field equations and their practical implications, which also apply to the classical Ridge penalty for linear models. We then embed our findings in the existing literature. Notably, we demonstrate that the Rudin-Osher-Fatemi model for image denoising can be recovered when the features are considered deterministic and evenly distributed.

The Vector Conditional Autoregressive Wishart Model for Multivariate Stock Market Volatility

Jan Vogler

The multivariate GARCH approach of Bollerslev, Engle and Wooldridge (1988), which in its vectorized formulation is known as the VEC model, is the most flexible multivariate time series model for conditional covariance matrices. However, estimation of VEC models based on daily returns of risky assets is computationally challenging so that their unrestricted estimation is feasible only for settings with a quite small number of assets. In this paper we suggest a novel VEC Conditional Autoregressive Wishart (VEC-CAW) model which is based on daily realized covariance matrices computed from high-frequency intraday returns. We analyze properties of this VEC-CAW and focus on its estimation. In order to make maximum likelihood (ML) estimation of the realized VEC-CAW model computationally feasible, we derive the analytical expression for the gradient of the log-likelihood under the conditional Wishart assumption for realized covariance matrices. Constrained optimization guaranteeing stationarity and positive definiteness proceeds using Bregman divergences. Doing so we successfully conduct unrestricted ML estimation of the VEC-CAW model with reasonable computation time. Then finite sample properties of the ML estimator are investigated in a Monte Carlo study with up to 10 assets. Further, we estimate an unrestricted VEC-CAW in an empirical illustration.

Regularized Online Distributional Regression

Florian Ziel

Large-scale streaming data are common in modern machine learning applications and have led to the development of online learning algorithms. Many fields, such as supply chain management, weather and meteorology, energy markets, and finance, have pivoted towards using probabilistic forecasts. This results in the need not only for accurate learning of the expected value but also for learning the conditional heteroskedasticity and conditional moments. Against this backdrop, we present a methodology for online estimation of regularized, linear distributional models. The proposed algorithm is based on a combination of recent developments for the online estimation of LASSO models and the well-known GAMLSS framework. We provide a case study on day-ahead electricity price forecasting, in which we show the competitive performance of the incremental estimation combined with strongly reduced computational effort. Our algorithms are implemented in a computationally efficient Python package ondil.

Learning of Smoothing Parameters in GAM Frameworks

Monika Zimmermann

Generalized Additive Models (GAMs) and their hybrid extensions with orthogonal neural network components appealingly combine interpretability with flexibility. In many modern applications, these models must be fitted to large data and parameter spaces. Existing strategies achieve scalability by exploiting sparsity - through banded Cholesky factorizations or reordering schemes that minimize fill-in - or by discretizing the model matrix. In both cases, the Cholesky decomposition remains the computational core for solving P-IRLS iterations and evaluating marginal-likelihood gradients in smoothing parameter updates. This work aims to enhance scalability by avoiding either explicit inversion or the entire Cholesky decomposition in favor of iterative, matrix-free techniques inspired by recent advances in the Gaussian process literature. When sparsity is unavailable and discretization insufficient, we propose to use Preconditioned Conjugate Gradient (PCG) methods—also termed Hessian-free optimization—combined with Hutchinson-type trace estimation within generalized Fellner–Schall iterations for model and smoothing parameter updates. Where Cholesky remains feasible, Hutchinson estimators still accelerate updates by replacing inversion with backward and forward solves. Initial results demonstrate superior scalability to parallel Cholesky methods in the non-sparse case, with stability ensured by low-rank pivoted Cholesky preconditioning for extreme smoothing parameters, consistent with Gaussian process findings. Future work will extend the approach to orthogonal hybrid GAM–NN models via K-FAC or PCG–Hutchinson Fellner–Schall type updates.

5. Abstracts of Invited Talks

Integrating Complex Covariate Transformations in Generalized Additive Models

Claudia Collarin

Transforming independent variables is common in applied statistics for improving model interpretability and ensuring assumptions for valid inference. While these transformations modify existing variables, feature engineering also creates new variables and uses dimension reduction, prioritizing predictive performance. Usually, both transformations and feature engineering occur during pre-processing. This project aims to fully integrate part of the feature engineering process into the modelling phase. We focus particularly on interpretable transformations designed to handle complex covariates, such as time series and spatial data, and on their embedding into multi-parameter GAMs. In particular, we extend GAMs to accommodate smooth effects whose argument is any scalar-valued, nested covariate transformation $\tilde{s} : \mathbb{R}^p \to \mathbb{R}$, that is fourth-order differentiable with respect to its own parameters. These transformations are treated as integral components of the model and their parameters are estimated jointly with the regression and smoothing coefficients. The gamFactory R package is implemented to incorporate this class of nested effects. It is integrated with mgcv, ensuring both a stable estimation process and a general application framework. To provide an example of application, we consider the problem of forecasting the total hourly net-demand, y_t, in Great Britain using, among other covariates, an hourly forecast of external temperature, temp_t. Due to the thermal inertia of the buildings, the consumption of electrical heating and cooling at time t is not entirely driven by temp_t, but depends on temp_{t-1}, temp_{t-2}, ..., as well. One way to capture thermal inertia is to include in the forecasting model an exponentially smoothed temperature covariate, temp_t^S.

Spatio-Temporal Count Autoregression

Kostas Fokianos

We study the problem of modeling and inference for spatio-temporal count processes. Our approach uses parsimonious parameterisations of multivariate autoregressive count time series models, including possible regression on covariates. We control the number of parameters by specifying spatial neighbourhood structures for possibly huge matrices that take into account spatio- temporal dependencies. This work is motivated by real data applications which call for suitable models. Extensive simulation studies show that our approach yields reliable estimators. This is a joint work with S. Maletz and R. Fried.

D-Vine Copula based Probabilistic Weather Forecasting

Annette Möller

Today's state-of-the-art approach to weather forecasting utilizes numerical weather prediction (NWP) models, which are physical models describing the dynamics of the atmosphere. Typically, an NWP model is run multiple times, with different initial conditions and/or model formulations to obtain a so-called forecast ensemble that represents model and forecast uncertainty. However, the resulting ensemble forecasts still tend to exhibit systematic bias and dispersion errors and consequently lack calibration, which means the ensemble is not able to correctly represent forecast uncertainty. Therefore, ensembles of NWP forecasts require correction by so-called postprocessing models. These approaches are based on statistical or machine learning models with the aim to obtain better calibrated probabilistic forecasts. We propose a D-vine copula based quantile regression (DVQR) approach for postprocessing of ensemble weather forecasts. It allows to choose important predictor variables for the D-vine regression by a sequential forward selection procedure from a potentially large set of variables, and is thus able to deal with high-dimensional predictor settings. The estimation algorithm is highly data driven and adopts more general dependence structures than the state-of-the-art Ensemble model output statistic (EMOS) postprocessing model. The current DVQR does not explicitly allow to account for e.g. temporal or spatio-temporal covariate effects. However, such effects can be expected to be present in weather data, specifically for variables such as temperature showing strong seasonal behaviour. We introduce an extension, where we parametrize the bivariate copulas in the D-vine through Kendall's tau, which can then be linked to additional covariates introducing temporal or spatial effects. The parametrization of the correlation parameter allows e.g. generalized additive models (GAMs) to detect potentially hidden covariate effects. The new method is called GAM-DVQR, and its performance is illustrated in a case study for postprocessing of 2m surface temperature forecasts.

High-Dimensional Moore-Penrose and Ridge-Type Regularization

Nestor Parolya

In high-dimensional statistics, generalized inverses of covariance matrices often appear implicitly in regularization and inference problems. In this talk, we investigate the asymptotic properties of the Moore-Penrose inverse and related ridge-type inverses of the sample covariance matrix for number of dimensions p greater than the sample size n. We derive analytical expressions for the limiting behavior of weighted trace functionals, represented via partial exponential Bell polynomials, which are both interpretable and computable in practice. The results go beyond existing theory by allowing for general population covariance structures, relaxing Gaussian assumptions, and operating under high-dimensional asymptotic regimes. From an applied perspective, these insights provide a universal methodology for constructing fully data-driven shrinkage estimators of precision matrices, with implications for spatio-temporal modeling, portfolio optimization, and other regularization-based inference tasks. The Moore-Penrose inverse itself emerges as a natural high-dimensional regularizer, with simple transformations yielding competitive or superior performance compared to classical benchmarks, while keeping computational costs low.

Ultra-High Dimensional Cointegration (with Shi Chen)

Melanie Schienle

This paper proposes a method for model determination in ultra-high dimensional cointegrated systems where the cross-section dimension m can even largely exceed the sample size T . For such ultra-high dimensional cases, we require an adequate non-standard prescreening step which we develop for the nonstationary cointegration vector but also for the stationary loading matrix. We prove that identified sets for the non-zero loadings and the cointegration space contain the respective true sets with high probability. A feasible algorithm is provided, making the technique easily accessible for practitioners. In a second step, we employ reduced rank regression based on the pre-selected set of variables, and show the cointegration rank selection consistency of the overall procedure. In order to achieve consistent rank selection, we propose a tailored information criterion which is also of general interest for factor models when both strong and weak factors are present. Results of the simulation study demonstrate competitive performance of the proposed methodology. In an empirical illustration with 1045 NASDAQ stocks, the proposed methodology allows for large-scale multivariate predictive regression for the entire system.

Sparse High-Dimensional Vector Autoregressive Bootstrap (with Robert Adamek and Ines Wilms)

Stephan Smeekes

We introduce a high-dimensional multiplier bootstrap for time series data based on capturing dependence through a sparsely estimated vector autoregressive model. We prove its consistency for inference on high-dimensional means under two different moment assumptions on the errors, namely sub-gaussian moments and a finite number of absolute moments. In establishing these results, we derive a Gaussian approximation for the maximum mean of a linear process, which may be of independent interest. Our simulation studies demonstrate the potential of the method in finite samples, even if the data are not generated by a finite-order vector autoregression.

Throwing Vines at the Wall: Structure Learning via Bias-Corrected Random Search

Thibault Vatter

Vine copulas, allowing to flexibly model complex multivariate dependencies, have become widely used in both statistics and machine learning, yet structure learning remains a key challenge. In that context, the standard greedy algorithm of Dissmann (2013) is efficient but often suboptimal. We propose randomized algorithms that improve structure selection and a statistical framework correcting for the optimism of validation-based model choice. The correction relies on Gaussian comparison inequalities, yielding theoretical guarantees on selection probabilities. Empirical results on synthetic and real-world data show that our methods outperform state-of-the-art approaches in both accuracy and reliability.

Inferring the spatio-temporal structure of climate variability and its external and internal drivers

Nils Weitzel

The climate system is a high-dimensional, complex system that is subject to fluctuations on timescales from minutes to millions of years. These are driven by changes in boundary conditions and internal dynamics. In this talk, I will first review statistical challenges in characterizing the spatio-temporal structure of climate variability, attributing climate change to internal and external sources, and incorporating constraints from past climate changes into future predictions. This includes methods for merging physical understanding and observations, dimension reduction, uncertainty quantification, and model selection. Then, I will present a Bayesian hierarchical framework for reconstructing spatio-temporal climate changes from natural climate archives, the only source of information on climate changes beyond the last 200 years. Incorporating this data requires the use of observational operators that link climate changes to measured quantities such as pollen assemblages and chemical element ratios. These operators can be non-linear and involve multiple sources of timescale-dependent uncertainties. Observational operators are also a key ingredient for evaluating the ability of Earth system simulators to reproduce past climate changes, which is an important out-of-sample test for simulators that predict future climate. Using spectral analysis and probabilistic evaluation methods, I deduce that simulators tend to have systematic deficiencies in the regional structure of climate variability despite accurately reproducing global mean variability. I will close my talk with some ideas for how to incorporate these insights into future predictions.