October 26, 2025

Title
Advances in Real-Time Processing of Longitudinal Data: From Statistical and Deep Learning Methods to Applications

Topics
Adaptive statistical filtering, Anomaly detection, Deep learning architectures, Forecasting, Multivariate time-series, Speech enhancement, Voice activity detection

Abstract
Longitudinal data are generated by many sensors that are being massively deployed across wide range of systems. This yields very large amounts of multi-dimensional and possibly heterogeneous time series data that must be processed to identify changes in the system behavior, to detect events, and to predict the future values in order to control the system states. The temporal dimension of longitudinal data prohibits processing all samples at once. It necessitates using statistical filtering techniques, especially when the processing needs to be performed in real-time, and when the outputs values need to be generated at the same rate as the input samples are arriving. The filtering must capture the important temporal patterns across many time-scales. Another challenge is how to effectively process multiple time-series to capture their spatial dependencies without excessively increasing the computational complexity. The traditional model-based approaches often rely on Gaussianity of samples, and their well-defined correlation patterns. However, in many scenarios, the samples are non-Gaussian, which requires adopting non-linear modeling techniques. When even non-linear models fail to achieve an acceptable accuracy, the sample-based universal model-free methods can be considered. These methods usually combine various signal decompositions with deep learning filtering. They are effective for whole classes of input signals rather than being tailored to one specific type of the signal unlike statistical filtering methods. The drawback of the model-free methods is that they often require time consuming and expensive training. They are also numerically much more complex, which can be prohibitive in some practical implementations. The objectives of this tutorial are as follows. (1) Describe the goals, challenges and opportunities in processing the real-world longitudinal data; (2) Explain the pros and cons of using traditional adaptive filtering techniques vs. using modern deep learning architectures; (3) Explore implementation constraints in developing practical applications assuming the speech applications as an example; and, (4) Summarize the recent literature, and outline the open research problems.

Draft schedule (3 hours tutorial)

Part 1: Adaptive statistical filtering for longitudinal data (1 hour)
Presenter: Ying-Ren Chien

- adaptive filtering: LMS, RLS, Kalman filter
- signal projections, kernel based methods, fractional filtering
- non-linear models: modified Kalman filter, Volterra and Koopman models
- other topics: coentropy, biases, sparsity, improving robustness, interference cancellation

Part 2: Deep learning methods for longitudinal data (1 hour)
Presenter: Pavel Loskot

- pre-processing multivariate time-series: normalizations and space-time decompositions
- predictions, anomaly detection in univariate time-series: gated and graph-based networks
- predicting multivariate time series: transformer-based and diffusion-based architectures
- hyper-parameters tuning, information leakage, regularizations, training strategies

Part 3: Implementation issues and applications
Presenter: Yu Gao

- embedding deep learning models in consumer electronics
- speech enhancement and speech extraction in noisy multi-talker scenarios
- sound separation, anomalous sound detection, sound events localization and detection
- voice activity detection and the keyword spotting, speech emotion recognition

Relevance to the conference

Many engineering systems are monitored by sensors to maintain their stable operation. This generates very large amounts of longitudinal data, which must be often processed in real-time. These data can be used for detecting events and changes in the system behavior as well as to build digital twin models of the systems. The underlying models may need to be continuously updated as the new observations are obtained. The proposed tutorial will survey the recent advances in processing longitudinal data, explain how the traditional approaches can be made more robust, and more widely applicable to different situations and scenarios. The recent trend in using deep learning architectures that are mostly based on transformers and diffusion will be explored. The last part of the tutorial will cover the implementation issues, and provides guidelines how the practical applications are being developed. The background of the tutorial presenters ensures that the theoretical results will be complemented with practical insights. One of the main goals is to explain when the traditional filtering methods may be sufficient, and when deep learning architectures are necessary, and how to make the latter to be effective for multivariate longitudinal data. The practical implementation constraints will also be considered.

Intended audience and level

The intended audience is graduate students and researchers with interest in modeling and processing time-series and other longitudinal data. The tutorial will be of interest to anyone who is looking for new research ideas in time series processing, and who is required to design and implement downstream applications that are powered by the longitudinal data. The tutorial will explain how to make traditional adaptive filtering more robust, and how to modify common deep learning architectures such as those using attention mechanisms or diffusion to be also effective for longitudinal data beyond natural language processing. The intended technical level of this tutorial is intermediate, i.e., it is expected that the attendees already have some knowledge in statistical signal processing and deep machine learning methods from their graduate courses. The first two tutorial presenters are experienced university professors with many years of research and teaching experience. The third tutorial speaker has in-depth understanding of practical issues with implementing signal processing and machine learning algorithms in consumer electronics.

Tutorial speakers

Prof. Ying-Ren Chien
National Taipei University of Technology, Taipei
yrchien@ntut.edu.tw
https://scholar.google.com/citations?user=Tm-IJfwAAAAJ

Biography: Ying-Ren Chien (Senior Member, IEEE) is a Full Professors in the Department of Electronic Engineering, National Taipei University of Technology, Taipei. His research interests include consumer electronics, multimedia denoising algorithms, adaptive signal processing theory, active noise control, machine learning, the Internet of Things, and interference cancellation. Dr. Chien was the recipient of the best paper awards, including ICCCAS 2007, ROCKLING 2017, IEEE ISPACS 2021, IEEE CESoc/CTSoc Service Awards in 2019, NSC/MOST Special Outstanding Talent Award in 2021, 2023, and 2024, Excellent Research-Teacher Award in 2018 and 2022, and Excellent Teaching Award in 2021. From 2023 to 2024, he was Vice Chair of the IEEE Consumer Technology Society Virtual Reality, Augmented Reality, and Metaverse (VAM) Technical Committee (TC). Since 2025, he has been the Secretary of IEEE CTSoc Audio/Video Systems and Signal Processing TC. He is currently an Associate Editor for IEEE Transactions on Consumer Electronics.

Assoc. Prof. Pavel Loskot
Zhejiang University – UIUC Institute, Haining
pavelloskot@intl.zju.edu.cn
https://scholar.google.com/citations?user=HOXVuYAAAAAJ

Biography: Pavel Loskot (Senior Member, IEEE) joined the ZJU-UIUC Institute in January 2021 as Associate Professor after 14 years with College of Engineering, Swansea University, UK. In the past 30 years, he was involved in numerous collaborative research and development projects, and also held a number of paid consultancy contracts with industry mainly, but not only in wireless communications. His research interests focus on mathematical and probabilistic modeling, statistical and digital signal processing, and machine learning for multi-sensor, tabular, and longitudinal data. He received 8 best paper awards, and delivered 18 tutorials in international conferences including **BigComp 2024**, APSIPA ASC 2017/2021/2022/2024, and IEEE MILCOM 2018/2019. He is the Fellow of the HEA, UK, Recognized Research Supervisor of the UKCGE, and the IARIA Fellow. He is the Editor in ICT Express.

Yu Gao
AI Research Center, Midea Group, Shanghai
gaoyu11@midea.com
https://scholar.google.com/citations?user=VnCFz24AAAAJ

Biography: Gao Yu obtained BSc and MSc degrees in EE from the USTC. He is currently the Head of Human-Computer Interaction Algorithms at Midea Group's AI Innovation Center, and the Director of the National New Generation AI Innovation Platform for Home Robots. He holds over 50 domestic and international patents in speech processing and NLP concerning intelligent speech & language algorithms, and the AI industrialization. He has led the development of 10+ IEEE and national standards. He is also an Executive Member of the CCF Speech & Dialogue Special Committee, and a Member of the National Standardization Committee (TC46/TC28).