Winter School in Network Data Science and Artificial Intelligence

9-13 February 2026
East Campus USI-SUPSI
Lugano-Viganello

USI-SUPSI PhD School
in Applied Data Science
and Artificial Intelligence

Monday, February 9th, 2026

12:00-13:00
Welcome Reception

13:00 - 16:00
Workshop
Introduction to Relational Event Models
Melania Lembo, Martina Boschi

16:00 - 16:30
Coffee Break

16:30 - 17:15
Keynote Lecture
Causal relational event models
Veronica Vinciotti

17:15 - 19:00
Apero
Open to all Winter School participants, instructors, and speakers

Tuesday, February 10th, 2026

09:00 - 09:45
Keynote Lecture
Assessing the Fit of Relational Event Models: A Simulation-based Approach Accounting for Temporal Structure
Viviana Amati

09:45 - 10:30
Keynote Lecture
Scalable Durational Event Models: Application to Physical and Digital Interactions
Cornelius Fritz

10:30 - 11:00
Coffee Break

11:00 - 11:45
Keynote Lecture
Not All Bonds Are Created Equal: Dyadic Latent Class Models for Relational Event Data
Joris Mulder

11:45 - 12:30
Keynote Lecture
Modeling Social Networks with Changeable Nodes
Alessandro Lomi

12:30 - 14:00
Lunch Break

14:00 - 16:45
Workshop
All You Need to Know About Relational Hyper-event Modeling
Martina Boschi, Melania Lembo

16:45 - 17:15
Coffee Break

17:15 - 18:00
Keynote Lecture
What and Whom Do We Cite? Relational Hyperevent Models for Citation Networks
Jürgen Lerner

Wednesday, February 11th, 2026

09:00 - 12:30
Social activity

12:30 - 14:00
Lunch

14:00 - 17:00
PhD students presentations

Thursday, February 12th, 2026

09:00 - 10:30
Tutorial
Graph Deep Learning for Time Series
and Spatiotemporal Data - part 1
Daniele Zambon

10:30 - 11:00
Coffee Break

11:00 - 12:30
Tutorial
Graph Deep Learning for Time Series
and Spatiotemporal Data - part 2
Daniele Zambon

12:30 - 14:00
Lunch Break

14:00 - 15:30
Tutorial
Learning dynamical systems
from data - part 1
Manas Mejari

15:30 - 16:00
Coffee Break

16:00 - 17:30
Tutorial
Learning dynamical systems
from data - part 2
Manas Mejari

Friday, February 13th, 2026

09:00 - 10:00
Keynote Lecture
Privacy-aware Probabilistic
Graphical Models
Cassio de Campos

10:00 - 11:00
Keynote Lecture
From Theory to Practice:
Overcoming the Real-World Challenges
of Reinforcement Learning
Marcello Restelli

11:00 - 11:30
Coffee Break

11:30- 12:30
Keynote Lecture
Building Context-Aware Foundation
Models for Time Series Forecasting
Abdul Fatir Ansari

12:30 - 14:00
Lunch Break and Closing

Abstracts

Workshop: Introduction to Relational Event Models

Melania Lembo, Martina Boschi

Abstract

Many sociological processes unfold as sequences of temporally ordered interaction events between entities. Leveraging the availability of time-stamped data, the Relational Event Model (REM) provides a powerful framework for generative modeling of such dynamic networks. REMs capture how node and edge characteristics, together with past interactions, shape the evolution of these processes. The essential temporal information embedded in relational event sequences has motivated advances in the specification of covariates and development of inference techniques.

In this workshop, we will introduce the foundations of REMs and show how to construct covariates that represent exogenous drivers, endogenous mechanisms, and temporal features of these processes.

We will then examine likelihood-based inference methods for estimating covariate effects and discuss extensions that enhance the flexibility of REMs, including non-linear and time-varying influences as well as random effects to account for network and actor heterogeneity.

Workshop Agenda

1. Core REMs: introduction to REM formulation and the inclusion of time-aware exogenous and endogenous variables.
2. REM inference: overview of likelihood-based estimation methods.
3. Non-Linear REMs: modeling non-linear and time-varying covariate effects while accounting for network and actor heterogeneity.
4. Practical Tutorial with R Studio: applying REMs to empirical datasets.

Keynote Lecture: Causal relational event models

Veronica Vinciotti

Abstract

Relational event models (REMs) describe temporal interactions between social actors, from the invasions of regions by alien species to financial transactions between banks or patient transfers between hospitals. An important question in all of these, is what are the causal drivers underlying these processes. This is a causal discovery question and requires a causal inferential

procedure that goes beyond MLE-based associative inference currently employed for REMs. We embed a causal REM within a structural causal framework by means of local conditional independence and causal validity for a definition of causality. By exploiting the connection between the REM partial likelihood under nested-case control sampling and logistic regression, we are able to use causal discovery methods recently developed for generalized linear models to identify the causal drivers of the relational event process. In particular, we show that a causal relational event model is identified by means of two conditions: a standard MLE property and a crucial invariance property, expressed in terms of Pearson risk. The empirical analogue of these conditions requires a statistical test for the detection of the causal model among a set of candidate models. We apply this method on a dataset of 350,000 events to investigate the causal drivers of bike sharing in Washington D.C.

Keynote Lecture: Assessing the Fit of Relational Event Models: A Simulation-based Approach Accounting for Temporal Structure

Viviana Amati

Abstract
We present a simulation-based procedure for assessing goodness of fit in relational event models. The approach extends the auxiliary variable method proposed in the literature on network modeling and compares statistics from observed relational event sequences with those from sequences simulated under the fitted model. A key contribution is the emphasis on the internal temporal structure of relational mechanisms, an aspect of model fit that has received limited attention in prior work. An application to patient-referral collaboration data among healthcare organizations demonstrates the ability of the proposed method to evaluate model fit.

Keynote Lecture: Scalable Durational Event Models: Application to Physical and Digital Interactions

Cornelius Fritz

Abstract
Durable interactions are ubiquitous in social network analysis and are increasingly observed with precise time stamps. Phone and video calls, for example, are events to which a specific duration can be assigned. We refer to this type of data encoding the start and end times of interactions as durational event data'. Recent advances in data collection have enabled the observation of such data over extended periods of time and between large populations of actors. Building on Relational Event Models, we propose the Durational Event Model as a framework for studying durational events by separately modeling event incidence and duration. To accommodate large-scale applications, we introduce a fast,

memory-efficient, and exact block-coordinate ascent algorithm. Theoretical and numerical results demonstrate several advantages of this approach over traditional Newton-Raphson-based methods. We apply the model to physical and digital interactions among college students in Copenhagen. Our empirical findings reveal that past interactions are the main drivers of physical interactions, whereas digital interactions are more strongly influenced by friendship ties and prior dyadic contact.

Keynote Lecture: Not All Bonds Are Created Equal: Dyadic Latent Class Models for Relational Event Data

Joris Mulder

Abstract

Dynamic social networks can be represented as time-ordered sequences of interactions between pairs of individuals. The relational event model (REM) is a central tool for analyzing such data, and existing approaches to unobserved heterogeneity typically rely on actor-level latent blocks. Under this framework, any latent variation in a REM coefficient has to be expressed through the latent classes assigned to the sender and receiver, because the model allows effects to vary only across those actor-level groupings. Yet many social processes generate pair-specific tendencies that cannot be decomposed into individual characteristics. Examples include persistent cooperation or conflict between particular country pairs, longstanding friendships or hostilities between individuals, or dyads with idiosyncratic histories. To capture this genuinely dyadic structure, we introduce a dyadic latent class relational event model (DLC-REM) in which effects depend on a dyad's latent class. This more flexible parameterization allows unobserved heterogeneity to be represented where it actually occurs: at the level of the dyad. Through simulations, we show that the DLC-REM can capture a broader range of data-generating processes than actor-based latent models—often with considerably fewer parameters. We illustrate the methodology using both simulated examples and an empirical application to relational event data.

Keynote Lecture: Modeling Social Networks with Changeable Nodes

Alessandro Lomi

Abstract

One of the most distinctive features of networks generated by the interdependent actions of social agents is the ability of network nodes tochange qualitative aspects of their internal composition—or, in other words, their identity. When nodes are able to reconfigure their internal structure over time, common attachment mechanisms

such as homophily become problematic and can no longer be assumed. This limitation is particularly stringent when nodes represent composite (or collective) agents, such as formal organizations or other corporate entities.

Ignoring what is arguably the most defining property of social agents as network nodes restricts our ability to understand how changes in the identity of the nodes and changes in network structure influence one another and coevolve over time. Against this general backdrop, this paper introduces a class of Stochastic Actor-Oriented Models (SAOMs) for mixed-mode networks that explicitly incorporates changes in internal structures as a core feature of network nodes. The model's primary objective is to capture the coevolution of the internal structure of network nodes and the external structure of the network connecting them.

Relaxing the assumption of fixed node identity creates a need to reconsider the notion of network autocorrelation. We therefore propose a new implementation of a decomposition method for network autocorrelation that accommodates shifts in the internal composition of nodes. Using data on collaborative relations among healthcare organizations, we demonstrate the model's capacity to reproduce the observed interorganizational network structure, the internal organizational structures of nodes, and field-level distributions of organizational activities and resources. We conclude by discussing how the proposed analytical approach can be extended to a broader variety of empirical settings where network nodes undergo observable changes in their internal structure.

Workshop: All You Need to Know About Relational Hyper-event Modeling

Martina Boschi, Melania Lembo

Abstract

Social network analysis has relied mainly on time-aggregated data and individual node connections. However, advances in data collection technologies now provide large amounts of time-stamped and hypernetwork information. Examples of such datasets include email multicast communication data and citation networks. Whereas methodological developments in temporal network science have expanded the toolkit for describing time-stamped and hypergraph data, relational hyper-event modeling has emerged as a central approach for the statistical modeling of the dynamic structures these data reveal. Recent contributions have made the Relational Hyper Event Model (RHEM) even more flexible, allowing it to address a range of applied demands.

In this workshop, we aim to present RHEMs, beginning with their core properties and progressing to their more novel modeling capacities, highlighting how they can be employed to address empirical questions.

Workshop Agenda

1. Core RHEMs: computing hyperevent covariates; formulating, estimating, and interpreting linear RHEMs.
2. Non-Linear RHEMs: modeling non-linear, time-varying covariate effects while accounting for network and actor heterogeneity.
3. Practical Tutorial with eventnet and R Studio: applying RHEMs to empirical datasets.

Keynote Lecture: What and whom do we cite? Relational hyperevent models for citation networks

Jürgen Lerner

Abstract

Scientific networks are often used to quantify science, ranging from the impact of researchers, journals, or universities over to measuring the interdisciplinarity or disruptiveness of scientific work. In this talk I present relational hyperevent models (RHEM) as a general model for the joint dynamics of coauthoring and citation networks. RHEM can account for polyadic dependencies in the data, i.e., interactions among more than just two nodes, which cannot be well represented by a purely dyadic specification of the event rate. Moreover, I discuss an extension of the model to RHEM with random node-level effects representing the latent popularity of papers and researchers and briefly overview practical aspects, including computational issues, in specifying and estimating RHEM with the open-source software eventnet (https://github.com/juergenlerner/eventnet).

Tutorial: Graph Deep Learning for Time Series and Spatiotemporal Data

Daniele Zambon

Abstract

In many application domains, including smart cities, environmental monitoring, and finance, data are generated across time and space by interrelated entities. The resulting spatiotemporal data form collections of time series that correlate due to spatial proximity, shared influences, or broader functional dependencies among the associated entities. While deep learning has proven effective for modeling complex temporal patterns, Graph Deep Learning (GDL) extends these capabilities by providing powerful and scalable tools for learning from spatially related data. Unlike traditional multivariate methods, GDL models such as graph neural networks (GNNs)

leverage pairwise relationships by conditioning predictions on generic and possibly dynamic graphs spanning the time series collection.
This tutorial introduces a methodological framework that unifies GDL with time series processing, supporting key tasks such as forecasting, reconstruction, and anomaly detection. The first part presents the modeling principles underlying the framework, reviews the main graph-based architectures commonly known as spatiotemporal graph neural networks (STGNNs), and outlines a general recipe for designing new models. The second part discusses challenges and emerging perspectives motivated by recent advances, with a focus on issues such as handling irregular data and learning relational information directly from observations.

Tutorial Agenda

1. The data: spatiotemporal data, correlated time series, and their graph-based representation.
2. GDL for time series processing: GNNs, STGNNs, and a general recipe for graph-based spatiotemporal modeling.
3. Challenges and tools: Dealing with irregular and missing data, inferring relational information from data, model quality assessment, open-source libraries, and future directions.

Tutorial: Learning dynamical systems from data

Manas Mejari

Abstract

A dynamical system is an object (or collection of objects) that evolves over time, possibly under external excitations. Dynamical systems are ubiquitous in nature and technology and arise in a variety of applications ranging from autonomous vehicles and robotics to process control, economics, and biological systems.
Learning dynamical models from data, traditionally known as System Identification, is a discipline that provides a principled framework for constructing mathematical models of system dynamics from data, when first-principles descriptions are incomplete or unavailable.
This workshop introduces core concepts in system identification, including model structures, recursive estimation techniques, and prediction-error methods. We will also cover the theoretical foundations of the classical Kalman filter, a fundamental tool for state estimation whose influence spans a wide range of engineering and scientific applications.

Tutorial Agenda

1. Dynamical systems: Introduction to dynamical systems and model structures.

2. System identification: Overview of system identification algorithms including recursive least-squares and prediction-error methods.
3. Filtering: Introduction to state-space representations and Kalman filtering.
4. Case studies: Demonstrations of identification and filtering techniques using Python-based simulations

Keynote Lecture: Privacy-aware Probabilistic Graphical Models

Cassio de Campos

Abstract

Probabilistic graphical models enable meaningful knowledge representation and reasoning. They have been proved effective across diverse domains, including healthcare, bioinformatics, economics, law, and image processing. As privacy concerns escalate, it becomes increasingly critical for publicly released models to safeguard sensitive information about the training data on which they were learned, but typically released models do not prioritise privacy by design. The main current practical approach involves introducing noise into the model's parameters. While this idea can protect against attacks, it also significantly impacts the model's usefulness. We present and discuss credal models as a practical approach for balancing privacy and utility. Credal models represent a set of models by allowing set-based parameter specifications. We will argue about ways that credal models can disguise the original model, therefore reducing the probability of successful attacks, while achieving meaningful inferential results. These discussion will be supported by some recent formal results and intuitive experiments illustrating the direction and how it compares to standard ideas, in particular the obfuscation by noise.

Keynote Lecture: Building Context-Aware Foundation Models for Time Series Forecasting

Abdul Fatir Ansari

Abstract

Time series foundation models (TSFMs) have brought a paradigm shift to forecasting. Unlike traditional models trained separately on individual datasets, TSFMs are pretrained once on large-scale time series data and then applied to a wide range of forecasting tasks in a "zero-shot" manner. However, a key limitation has hindered their broader adoption in production: most pretrained models handle only univariate data, relying solely on the historical observations of a single series. In real-world applications, exogenous contextual factors are often essential for accuracy: retail demand depends on promotions, while energy consumption

varies with weather. Moreover, many domains involve multiple co-evolving series that must be forecast jointly, such as system metrics like CPU usage, memory, and storage I/O. In this talk, I will introduce Chronos-2, a pretrained model designed to tackle arbitrary forecasting tasks — univariate, multivariate, and covariate-informed — in a zero-shot fashion. Chronos-2 uses in-context learning (ICL) to incorporate time-varying contextual signals, enabling effective multivariate and covariate-based forecasting. Its enhanced ICL mechanisms also improve univariate forecasting through cross learning, where information is shared across univariate series within a batch, boosting prediction accuracy. Chronos-2 marks a major advancement over prior TSFMs, achieving state-of-the-art performance across benchmarks and delivering substantial gains on covariate-informed tasks. By addressing core limitations of existing TSFMs, it establishes itself as a practical, general-purpose forecasting model ready for production deployment.

Keynote Lecture: From Theory to Practice: Overcoming the Real-World Challenges of Reinforcement

Marcello Restelli

Abstract

Reinforcement learning (RL) has achieved remarkable progress in simulated environments, yet deploying RL in real-world systems remains profoundly challenging. This lecture will explore the key obstacles that arise when moving from controlled benchmarks to dynamic, high-stakes applications such as robotics, healthcare, and industrial decision-making. We will examine issues related to sample inefficiency, safety and risk sensitivity, reward design, distributional shift, partial observability, and the difficulty of integrating prior knowledge. The talk will also highlight practical concerns, including unreliable simulators, expensive data collection, non-stationary environments, hardware constraints, and the need for human oversight. Finally, we will discuss emerging solutions—such as offline and batch RL, model-based approaches, safe RL frameworks, and hybrid learning-control architectures—that aim to bridge the gap between theory and deployment. Attendees will gain a clear understanding of why real-world RL is hard, what progress is being made, and where impactful research opportunities remain.