

CONTRIBUTED TALKS: AGENDA

NOV 26, 12.00-12.20: LEARNING DYNAMICAL SYSTEMS FROM DATA: A SIMPLE CROSS-VALIDATION PERSPECTIVE

BOUMEDIENE HAMZI AND HOUMAN OWHADI

Regressing the vector field of a dynamical system from a finite number of observed states is a natural way to learn surrogate models for such systems. We present variants of cross-validation (Kernel Flows and its variants based on Maximum Mean Discrepancy and Lyapunov exponents) as simple approaches for learning the kernel used in these emulators. More elements are at <https://arxiv.org/abs/2007.05074>

NOV 26, 12.20-12.40: A NEURAL NETWORK ENSEMBLE APPROACH FOR GDP FORECASTING
LUIGI LONGO, MASSIMO RICCABONI AND ARMANDO RUNGI

We propose an ensemble learning approach to forecast the future US GDP growth release. Our approach combines a Recurrent Neural Network (RNN) with a Dynamic Factor model accounting for time-variation in mean with a Generalized Autoregressive Score (DFM-GAS). The analysis is based on a dataset of predictors encompassing a wide range of variables measured at different frequencies. The forecast exercise aims at evaluating the ability of each model in the prediction of GDP dynamics accounting for breaks in mean, potentially caused by crisis and recessions affecting the economy. We analyze the weights of the ensemble model to show whether neural networks improve the forecast performance of the GDP growth rate in the aftermath of the global financial crisis starting in 2008, and at the beginning of the more recent Covid-19 crisis.

NOV 26, 12.40-13.0: DEEP LEARNING MODELING OF THE LIMIT ORDER BOOK: A COMPARATIVE PERSPECTIVE

JEREMY TURIEL, ANTONIO BRIOLA AND TOMASO ASTE

We address theoretical and practical questions in the domain of Deep Learning for High Frequency Trading. State-of-the-art models such as Random models, Logistic Regressions, LSTMs, LSTMs equipped with an Attention mask, CNN-LSTMs and MLPs are reviewed and compared on the same tasks, feature space, and dataset and clustered according to pairwise similarity and performance metrics. The underlying dimensions of the modelling techniques are hence investigated to understand whether these are intrinsic to the Limit Order Book's dynamics. We observe that the Multilayer Perceptron performs comparably to or better than state-of-the-art CNN-LSTM architectures indicating that dynamic spatial and temporal dimensions are a good approximation of the LOB's dynamics, but not necessarily the true underlying dimensions.

NOV 26, 13.00-13.20: EFFICIENT LEARNING IN RECURRENT NEURAL NETWORKS
DANIELE DI SARLI, CLAUDIO GALLICCHIO AND ALESSIO MICHELI

The paradigm of Reservoir Computing (RC) has proven to be a powerful yet extremely efficient approach for designing and training Recurrent Neural Networks (RNNs). The widely known Echo State Network is a RC model that inherits the same architecture of a vanilla RNN, but thanks to a proper initialization of the parameters in the state transition function it allows to avoid any kind of training for the recurrent neurons in the network by exploiting properties of the underlying dynamical system. Unfortunately, this approach does not perform well on data exhibiting long-term dependencies. Can this limitation be overcome?

NOV 27, 12.00-12.20: UNIT ROOT TEST COMBINATION VIA RANDOM FORESTS
DANIEL OLLECH, LUCA NOCCIOLA AND KARSTEN WEBEL

A large number of non-seasonal and seasonal unit root tests are available. Yet, it can be ambiguous which tests to rely on due to uncertainties in the data generating process. We evaluate the accuracy of a large set of unit root tests on time series that are simulated to be representative for economic time series in the M4 competition data. Further, using a conditional random forest based elimination algorithm, we assess which tests should be combined to improve on the accuracy of the single tests.

NOV 27, 12.20-12.40: A KERNEL TWO-SAMPLE TEST FOR FUNCTIONAL DATA
GEORGE WYNNE AND ANDREW DUNCAN

Kernel two-sample tests have found wide success in machine learning in the last 10 years due to their ability to perform computationally efficient non-parametric two-sample tests for which ever data sources one can define a kernel over. For example they have been used to train GANs and perform inference for graph data. Until now little attention has been made on functional data, where each data point is a discretised function, which is of particular importance to dynamical systems. We propose kernels for such data, extending the existing theory for data living in a finite dimensional space to an infinite dimensional space and show numerical evidence supporting the proposal against existing methods. Joint work with Andrew B. Duncan. Pre-print link: <https://arxiv.org/pdf/2008.11095.pdf>

NOV 27, 12.40-13.00: DYNAMIC TRADING WITH REINFORCEMENT LEARNING
ALESSIO BRINI AND DANIELE TANTARI

Portfolio optimization is an active field of research in the financial domain, that is abundant of theoretical and practical applications to build profitable trading systems. Beyond the hype surrounding the application of machine learning techniques to financial trading, reinforcement learning (RL) represents a framework for optimizing sequential decisions over the course of time. It is of interest to see if this set of algorithms can outperform standard optimal control techniques as dynamic programming and achieve meaningful trading signals. We generated a series of synthetic experiments to assess the ability of standard RL algorithms to retrieve the underlying known dynamics of the data. Results are compared to those obtained by the closed-form optimal solution for the considered problem. In particular, RL with function approximators allows to extend the financial environment towards a more complex structure with a richer state representation that can be applied also to real data and solve the problem in a data-driven manner.