Course Form for PKU Summer School International 2018

Course Title	**Machine Learning for Time Series Analysis – Statistical Models and Deep Learning** **机器学习与时间序列分析**
Teacher	LIU Yan
First day of classes	July 23, 2018
Last day of classes	July 27, 2018
Course Credit	2 credits

Course Description

Objective

The course aims to introduce machine learning models, including both statistical models and deep learning, to students in various science disciplines such as computer science, statistics, economics, finance, electronic engineering, biology, physics etc., who are interested in machine learning and statistical models for time series analysis and forecasting. Enrolled students should have basic knowledge in statistics and probability, linear algebra, optimization. No programming skills are required, but could be helpful for hands-on exercise. The class covers popular time series models, including vector-regressive models (VAR), ARIMA models, hidden Markov models, Karman filtering, as well as advanced models, such as neural network models (Long Short-term memory neural networks, recurrent neural networks, gated recurrent neural networks), support vector machine regression, Hawkes processes, sparse VAR models etc. At the end of the course, the students are expected to be able to do the following: (1) understanding the mathematical formulation of time series models; (2) apply time series models to real-application data; (3) potential of developing novel machine learning models for time series applications for publications.

Pre-requisites /Target audience

Pre-requisites: basic knowledge in statistics and probability, linear algebra, optimization. No programming skills are required, but could be helpful for hands-on exercise.
Target audience: Senior undergraduate students and graduate students in various disciplines (computer science, statistics, economics, finance, electronic engineering, biology, physics)

Proceeding of the Course

No

Assignments (essay or other forms)

Readings, In-class and online discussions, and take-home exercises

Evaluation Details

Attendance and Discussions: 25%
Assignments: 45%
Exam: 30%

Text Books and Reading Materials

- James D. Hamilton (1994). Time Series Analysis. Princeton Press;
- Ian Goodfellow and Yoshua Bengio and Aaron Courville (2016). Deep Learning, MIT Press.
- Additional readings are given in the Class Schedule.

Academic Integrity (If necessary)

Students are allowed to discuss readings and assignments among classmates in and outside the class. However, individual-based writing assignments must be independently completed (i.e., without any plagiarism).

CLASS SCHEDULE
(Subject to adjustment)

Session 1: Introduction to Time Series	Date: 7/23/2018

【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.)
Definition of time series, stationary and non-stationary time series, white noise
Applications of time series analysis and forecasting
Introduction of basic time series models, moving average, auto-regression, ARMA models and extensions

【Questions】
What is time series? What is stationary time series and non-stationary time series? What is the basic models for time series analysis

【Readings, Websites or Video Clips】
1. Hamilton Ch 1-4, 11

【Assignments for this session (if any)】
Exercise 1. Questions on basic definitions of time series and stationary time series
Exercise 2. Hands-on exercise on running R-code for moving average, auto-regression, ARMA models

Session 2: State-space models	Date: 7/24/2018

【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.)
Introduction of state-space models, including hidden Markov model and Kalman filter

【Questions】
What is state-space models? What is hidden Markov model and Kalma filter? When should we apply these models

【Readings, Websites or Video Clips】

1. Hamilton Ch 5
2. Reading- Rabiner 1986

【Assignments for this session (if any)】
Exercise 3. Questions on hidden Markov models and Kalman filtering
Exercise 4. Hands-on exercise on hidden Markov models applying to text modeling

Session 3: Neural Network models for time series	Date: 7/25/2018

【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.)
Introduction of basic neural network models, recurrent neural networks (RNN), long short-term memory neural networks, gated recurrent neural networks

【Questions】
What is RNN? What type of properties does RNN capture in time series? How to apply RNN for time series forecasting and prediction

【Readings, Websites or Video Clips】
1. Goodfellow et al, 2016 Chapter 4-6, 10

【Assignments for this session (if any)】
Exercise 5: Hands-on exercise on RNN for time series analysis

Session 4: Sparse VAR models and Granger causality	Date: 7/26/2018

【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.)
Introduction of lasso, VAR models, sparse VAR models, Granger causality and extensions

【Questions】
How can we learn temporal dependencies? How can we address problems in practical applications, such as nonstationary, irregular time series, relational time series?

【Readings, Websites or Video Clips】
1. Reading- Liu 2018

【Assignments for this session (if any)】
Exercise 6: Hands-on exercise on sparse-VAR models for time series dependence analysis

Session 5: Hawkes Process, and Support Vector Regression	Date: 7/27/2018

【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.)
Introduction of Poisson process, Hawkes processes, support vector machines, and support vector regression

【Questions】
How can we model stochastic data? How can we go beyond linear predictors?

【Readings, Websites or Video Clips】
1. Reading- Laub et al, 2015
2. Reading- Burges, 1998

【Assignments for this session (if any)】
Exercise 7. Hands-on exercise on Hawkes process models

Sources of Textbooks, Videos, and Additional Readings
Textbooks/Videos: • N/A
Additional Readings: • Rabiner. An introduction of Hidden Markov Model, 1986 • Patrick J. Laub, Thomas Taimre, Philip K. Pollett. Hawkes Processes. 2015. • Yan Liu. A tutorial on sparse vector autoregression and granger causality, 2018. • Christopher J.C. Burges. A Tutorial on Support Vector Machines for Pattern Recognition.