Course Form for PKU Summer School International 2019

Course Title	**Modern Machine Learning in Practice**
	当代机器学习实践
Teacher	YANG Zhirong
First day of classes	July 1, 2019
Last day of classes	July 12, 2019
Course Credit	3 credits

Course Description

Objective

AI, enabled by machine learning (ML) technology, has become the new electricity. The next-generation industry demands a wide range of AI or ML applications. However, when trying machine learning to a real-world problem, you probably find enormous difference from what you have learned from textbook or conventional courses which only focuses on theories and models. This course aims to bridge the gap and to help you develop realistic machine learning products. After the course, you will understand a set of common problems in practice as well as their related concepts, ranging from preprocessing, inference, diagnosis, and interpretation. In solutions, we carefully select one or two state-of-the-art methods for each problem, instead of overwhelming you by a tedious list of alternatives. In the course we also discuss some open problems and their possible solutions in frontier, so that you can pay attention and seek answers even after the course. The course assignment includes a small project for practice, where you will be guided in building a real ML solution starting from raw data.

Pre-requisites /Target audience

Data Structures and Algorithm, Probability and Mathematical Statistics.
Undergraduate students.

Proceeding of the Course

None

Assignments (essay or other forms)

Reading, Assignment and Programming

Evaluation Details

Attendance and Reading: 30%
Programming Project: 40%
Presentation: 30%

Text Books and Reading Materials

Dipanjan Sarkar, Raghav Bali, Tushar Sharma. Practical Machine Learning with Python. Apress, 2017.

Sunila Gollapudi. Practical Machine Learning. Packt Publishing, 2016.
Ian Goodfellow, Yoshua Bengio, Aaron Courville. Deep Learning. MIT Press, 2016.
Christopher Bishop. Pattern Recognition and Machine Learning. Springer, 2013.
Trevor Hastie, Robert Tibshirani, Jerome Friedman. The Elements of Statistical Learning. Springer, 2016.
Alex Smola and S.V.N. Vishwanathan. Introduction to Machine Learning. Cambridge University Press, 2010.

Academic Integrity (If necessary)

CLASS SCHEDULE
(Subject to adjustment)

Session 1： Introduction	Date:7/1 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Introduce machine learning basics, including mathematical preliminaries, machine learning concepts, and a typical machine project procedure.**	
Session 2：Prediction models （part I)	Date:7/2 9-12 AM
Session 3：Prediction models （part II)	Date:7/3 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Introduce modern prediction models, with focus on ensemble learning methods.**	
Session 4：Feature engineering	Date:7/4 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Introduce preprocessing steps and techniques for preparing machine learning data to handle issues such as missing values, imbalanced cases, feature selection, active learning.**	
Practice & Project: Prediction Models	Date:7/4 1-5 PM
Session 5：Model diagnosis	Date:7/5 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Introduce model diagnosis techniques, especially how to use visualization tools for debugging**	
Practice and Project: Feature engineering	Date:7/5

	1-5 PM
Session 6：Model evaluation and interpretation（part I)	Date:7/8 9-12 AM
Session 7：Model evaluation and interpretation （part II)	Date:7/9 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Introduce concepts and methods for model evaluation and interpretation.**	
Practice & Project: visualization	Date:7/19 1-5 PM
Session 8：Case studies	Date:7/10 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Present several real-world machine learning case studies.**	
Practice & Project: interpretation	Date:7/10 1-5 PM
Session 9：Frontiers（part I)	Date:7/11 9-12 AM
Session 10：Frontiers （part II)	Date:7/12 9-12 AM
【Description of the Session】(purpose, requirements, class and presentations scheduling, etc.) **Introduce several open problems in machine learning frontiers, including automatic ML, reinforcement learning, transfer learning, ML with privacy protection, simulator-based ML, fairness in ML, ML in adversarial settings, etc.**	
Practice & Project: project presentations	Date:7/12 PM