

Statistical Methods for Intelligent Information Processing
(SMIIP)

Lecture 0:
Introduction to the Course

Shuigeng Zhou

School of Computer Science
September 13, 2017



What is this course?



- It is about statistical methods, but it is NOT statistics
- It is heavily overlapped with machine learning, and is closely related to data mining
- Mostly, it is, but NOT all, machine learning
- It is more like **machine learning** than data mining

What is this course?



This course

Data mining

Machine learning

Statistics

Instructor



- Prof. Shuigeng Zhou
 - R502, Yifu Building, Fudan Handan campus
 - Tel: 55664967
 - Email: sgzhou@fudan.edu.cn
 - Homepage: <http://admis.fudan.edu.cn/~sgzhou>
- The best way to reach me is email!

Class Time and Venue



■ Time

- 1:30PM – 3:45PM, Wednesdays from September 13, 2017 to January 3, 2018
- Maybe, there will be no class on October 4 due to the National Day holiday
- So possibly there will be totally 16 weeks of class

■ Venue: HY605

Prerequisites



- **Mathematics:** probability theory & statistics, linear algebra, optimization theory
- **Computer science:** programming, data structures and algorithms

Grading



- 5%: class participation
- 10%: paper presentation
- 60%: projects and reports (4 projects, one project per month)
- 25%: Exam

Course Goal



- Presenting some typical **statistical methods** for intelligent information processing
- Introducing some **applications** of these methods
- The students are required to finish some **course projects** so that they can understand these methods well
- **Learn less, learn better**




Major Topics

- Probability and Statistics (概率与统计基础)
- Naïve Bayes classification & regression (贝叶斯分类与回归)
- Expectation Maximization (期望最大算法)
- Hidden Markov Models (隐马尔科夫模型)
- Conditional Random Field (条件随机场)
- Topic Models (主题模型)
- Random Walk (随机游走)
- Random Projection (随机投影)

Recommended Books / Materials



- Kevin P. Murphy, *Machine Learning: a Probabilistic Perspective*, MIT Press, 2012
- C. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2007
- Charles Sutton and Andrew McCallum, *An Introduction to Conditional Random Fields*, 2012
- Gregory F. Lawler and Vlada Limic, *Random Walk: A Modern Introduction*, 2010
- Other readings will come from freely available online material



Journals & Conferences



■ Journals

- AI, JMLR
- IEEE TPAMI, IEEE TKDE

■ Conferences

- Artificial intelligence: **AAAI, IJCAI, UAI, KR**
- Neural information processing: **NIPS**
- Machine learning: **ICML, COLT, ECML**
- Computer vision: **CVPR, ICCV, ECCV**
- Data mining: **KDD**

Thanks!



Questions?