OR 750. Advanced Data Analytics

Department of Systems Engineering and Operations Research
George Mason University
Fall 2017

The class will consist of 7-8 lectures given by the instructor on several advanced topics in data analysis. The rest of the semester (another 6-7 lectures) students will present on the topic of their choice.

Announcements

- 5/3/2017: First class is on Aug 30 at 4:30pm

List of topics

- Probabilistic models for Machine Learning
  - Conjugate distributions, exponential family
  - Model choice
  - Hierarchical linear and generalize linear models (regression and classification)
  - Models for missing data (EM-algorithm)
  - LDA, Normal mixtures, Bayes PCA
  - Bayes computations (MCMC, Variational Bayes)
  - Graphical Models
  - Probabilistic modeling with Stan
- Deep learning
  - Optimization
  - Architectures (CNN, LSTM, MP, VAE)
  - Bayesian DL
  - Generative models (GANs)
  - Modeling with TensorFlow
- Filtering
  - Kalman Filter
  - Extended KF and ensemble KF
  - Particle filter
  - Modeling with DLM package in R
- Probability and Statistics methods for Decision Making

    - Real-time hypothesis testing
    - Brownian motion
    - Bayesian methods for optimal stopping time detection
- Bayesian Optimization
    - Tuning machine learning algorithms
    - Engineering model calibration
    - Modeling with spearmint

Course staff

Instructor: Vadim Sokolov
Office: Engineering Building, Room 2242
vsokolov(at)gmu.edu
Tel: 703-993-4533

Other Reading

- Lectures on “Probability in High Dimension” (pdf)
- Book on “High-Dimensional Probability” (pdf)
- Blog post on SGD (link)

Office hours

Vadim Sokolov: By appointment (at Engineering 2242)

Lectures

Location: Nguyen Engineering Building 1109
Times: 4:30-7:10pm on Wednesday

Grades

Grade composition: No in-class examination. Grade is based entirely on participation in class and homework assignments.

Deep Learning

- Normalization Propagation: A Parametric Technique for Removing Internal Covariate Shift in Deep Networks (paper)
- Why does Monte Carlo Fail to Work Properly in High-Dimensional Optimization Problems? (paper)
- Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift (paper)
- Auto-Encoding Variational Bayes (paper)

- A Central Limit Theorem for Convex Sets (paper, slides)
- Learning Deep Architectures for AI monograph
- Generative Adversarial Networks (presentation)
- GANs at OpenAI (blog)
- Tuning CNN architecture (blog)
- Unsupervised learning (blog)
- Deep Energy (blog)
- DL Summer school 2015 (videos)
- DL Representations (blog)

Decision Making

- Multivariate Industrial Time Series with Cyber-Attack Simulation: Fault Detection Using an LSTM-based Predictive Data Model (paper)

Filtering

- Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems paper
- Can local particle filters beat the curse of dimensionality? paper

Bayes

- The Markov Chain Monte Carlo Revolution (paper)
- Graphical Models, Exponential Families, and Variational Inference (monograph)
- Variational Inference: A Review for Statisticians (paper)

Calibration

- Hyperband (demo)

Tools

- Stan
- TensorFlow
- dlm (R)
- Keras

Misc Links

- Independent learning