

APAM 2610

Recent Applications of Probability and Statistics

I. Intuition and large systems

- A. Gibbs ensemble and the maximum-entropy principle
 - 1. The Gibbs ensemble
 - 2. The “maximum-entropy principle”
- B. Large deviations and information theory
 - 1. Gibbs ensemble with a prior probability distribution
 - 2. Connections to information theory
- C. Asymptotic independence
 - 1. Exchangeability
 - 2. Local chaos

II. Some statistical theory

- A. Estimation
 - 1. The estimation problem
 - 2. Consistency
 - 3. Bias and variance
 - 4. The maximum-likelihood principle
- B. Classification
 - 1. The classification problem
 - 2. Neyman-Pearson lemma and God’s ROC curve
- C. Building classifiers – nonparametrically
 - 1. k-nearest-neighbor classifier
 - 2. Support-vector machine
 - 3. Cross-validation
 - 4. Vapnik-Chervonenkis dimension and other measures of model complexity

III. Dependency graphs and computing

- A. Markov random fields and Gibbs distributions
 - 1. Definitions and equivalence
 - 2. Marginals and posteriors
- B. Computing
 - 1. Dynamic programming on general graphs
 - 2. MCMC on general graphs
 - 3. Example: application to image processing
 - 4. Coupling and exact sampling
- C. Estimation
 - 1. Exponential families
 - 2. Maximum likelihood
 - 3. Pseudo-likelihood
 - 4. MM and EM

APAM 2620 (proposed)

IV. Some Bayesian statistical theory

- A. The estimation problem
- B. Maximum likelihood and failure of maximum likelihood on “large” problems
- C. Penalized likelihood and prior probabilities
- D. Some general properties of Bayesian estimators
- E. Conjugate priors
- F. Dirichlet priors
- G. Applications to document classification and hierarchical Markov models

V. Filtering

- A. Discrete-state hidden Markov models and the break-down of dynamic programming with large state spaces
- B. Continuous state spaces and the Kalman filter
- C. Particle filters
- D. Applications to Black-Sholes theory, predicting the markets, and making money