GEORGIA INSTITUTE OF TECHNOLOGY

7751
Probabilistic Graphical Models in Machine Learning
Spring Semester 2027

Instructor:

Faramarz Fekri
email: fekri@ece.gatech.edu

Head TA:

TBD

HWK Graders

TBD

Class Hours:

MW: 2:0-3:15 PM (Location: Weber SST III)
All Lectures will be in person for Section A (campus students). Sections Q (off campus students) will use recorded video lectures of the class. Live zoom will be available too for Sections Q.

Office Hours:

Campus Students: MW after class for brief questions,
also Mondays 4:00PM-5:30PM, Coffee shop in Clough Commons

Off campus students: Mondays 9:00PM-10:00PM EST through zoom.

Course Delivery Mode:

Lectures will be recorded and posted on Canvas for all students access.

Attendance:

Regular attendance of in person lectures are **mandatory** for campus students, unless you have a valid reason to be absent. Occasional class attendees will be taken and will be used toward final grading.

Course Objectives:

The course will provide students with an introduction to the theory and practice of graphical models, one of the most dominant frameworks in machine learning and artificial intelligence. The class will cover three main aspects: The core **representation**: including Bayesian and Markov networks, and dynamic Bayesian networks; **probabilistic inference algorithms**: both exact and approximate methods; and **learning methods**: for both the parameters and the structure of graphical models. Specifically, students will:

1. Become familiar with the most commonly used graphical model representation methods, learning and inference algorithms.
2. Gain exposure to the application of graphical models to real world problems.

3. Learn as to how one can formulate a wide range of problems with very large number of variables using the unified language of graphical models.

Textbook:

- Most of the materials are covered in the book "Probabilistic Graphical Models: Principles and Techniques," Daphne Koller & Nir Friedman.
- Lecture Slides will be provided.
- A few topics will be based on the book "An Introduction to Probabilistic Graphical Models" in preparation by Michael I. Jordan. PDF of a few Chapters of the book (as the "duplicate notes") will be provided as the course progresses.
- Additional Reference: "Machine Learning A Probabilistic Perspective,'' Kevin P. Murphy, MIT Press.

Course Prerequisites:

Familiarity with basic linear algebra and very basic statistics and probability theory is assumed.

Grading Formula for Campus Students:

Four Homeworks (including mini-projects)	80%
Four Pop Quizzes	20%

Grading Formula for Section Q (off campus Students):

Four Homeworks (including mini-projects)	100%

Pop Quizzes:

There will be total of 4 pop quizzes for campus students, each having 5 point credit. Attending the quiz counts 1 point and having correct solution counts 4 points in each quiz.

Homework:

The primary way to learn any subject is to WORK HOMEWORK PROBLEMS: as many as possible, and work them CAREFULLY. There will be 4 Homeworks. They will be due at the beginning of the class on due dates (which are approximately three weeks from the dates they were assigned). Late homework will NOT be accepted for grading**.** Homework is to be written up and submitted (PDF and codes) online to canvas individually. Working with colleagues is encouraged but simply copying someone else's solution is not acceptable and will be treated as such. Homework will be graded and solutions will be available. All campus and off campus students will have the same due date.

Honor Code:

Academic dishonesty will not be tolerated. This includes cheating, lying about course matters, plagiarism, or helping others commit a violation of the Honor Code. Plagiarism includes reproducing the words of others without both the use of quotation marks and citation. Students are reminded of the obligations and

expectations associated with the Georgia Tech Academic Honor Code and Student Code of Conduct, available online at www.honor.gatech.edu.

Canvas:

All course materials (lecture videos, slides, HWKs, Solutions to HWKs, etc.) can be found at http://canvas.gatech.edu

Topical Outline:

1. Bayesian networks *(Representation)*

   a. Examples (HMM, diagnostic system, etc.)
   b. Separation and independence
   c. Markov properties and minimalism
   d. Applications to time series model, topic modeling and network modeling

2. Markov networks *(Representation)*

   a. Examples (Boltzmann machine, Markov random field, Ising models)
   b. Cliques and potentials
   c. Markov properties
   d. Factor graphs
   e. Applications to image modeling, and network modeling

3. Gaussian network models and exponential family

   a. Multivariate Gaussians and Gaussian Networks
   b. Exponential families
   c. Related topics from Information Theory

4. Exact inference *(Inference)*

   a. Complexity
   b. Variable elimination
   c. Belief propagation (message passing) on trees
   d. Sum- and Max-product algorithms
   e. Clique tree, Junction tree
   f. Application to HMM

5. Sampling methods for approximate inference *(Inference)*

   a. Particle Filtering
   b. Rejection Sampling Method
   c. Likelihood Weighting
   d. Importance Sampling Method
   e. MCMC Sampling

i. Gibbs Sampling Method
ii. Metropolis–Hastings method

6. More on approximate inference *(Inference)*

a. Loopy belief propagation
b. Variational inference and optimization view of inference
c. Mean field approach

7. Parameter learning (*Learning*)

a. Parameterizing graphical models
b. Parameter estimation in fully observed Bayesian networks
- Maximum likelihood estimation
- Bayesian parameter estimation
- Examples: HMM, etc.

c. Parameter estimation in fully observed Markov networks
- Maximum likelihood estimation
  - Iterative Proportional Fitting (IPF)
  - Generalized Iterative Scaling (GIS)

d. Parameter estimation in partially observed graphical models
- Expectation-Maximization (EM)
- Examples: HMM, etc.

e. Learning Conditional Random Fields

8. Structure learning (*Learning*)

a. Score based approach
b. Chow-Liu algorithm for Bayesian networks
c. l1-regularized convex optimization for Markov random fields
d. Low-rank regularized learning of latent variable models

9. Nonparametric Bayes methods (*Learning*) (time permitting)

a. Gaussian processes
b. Dirichlet processes
c. Indian Buffet processes