Brennan Klein (he/him)
b.klein@northeastern.edu
Network Science Institute
Northeastern University

NETS 7976: Directed Study
Reasoning Under Uncertainty with
Probabilistic Machine Learning

Fall 2023; 2 credits
Thursday: 12:00pm – 1:30pm
Sep 7 – Dec 13, 2023
177 Huntington, #1005

Summary

This class is an introduction to the field of *Probabilistic Machine Learning*. The course will build the mathematical foundation for understanding inference, learning, and reasoning through the lens of probability theory. This is especially relevant in today's scientific landscape, given that the theoretical backbone of modern machine learning and artificial intelligence techniques is inextricably built on fundamental concepts from probability theory. Students will leave the class with experience in solving theoretical and applied problems related to Probabilistic Machine Learning. This class is designed in collaboration with **Moritz Laber**, a second year PhD student in Network Science.

Coursework, Class Structure, Grading

The course adopts a flipped classroom concept: Students are exposed to the relevant material through pre-recorded video lectures, readings, and exercises that involve both analytic calculation and coding. The in-class hours are dedicated to recapitulation of the most important points, clarification of questions, as well as open discussion. The grading of the course is based on active participation in discussions and continual work on exercises. In a typical class students will summarize the main points of the lectures in their own words. The weekly exercises and questions, that students prepare prior to class, structure the in class discussion. Discussions will remain open and allow to explore topics deeper as needed.

Learning Objectives and Outcomes

By the end of this course students should have a deep familiarity with the central concepts of probabilistic machine learning and reasoning under uncertainty. In particular, they develop a thorough understanding of Gaussian processes (GP) for regression and classification, from mathematical, algorithmic and applied perspectives. Building on this understanding, they learn about the relevance of uncertainty for deep learning and how this problems can be addressed through the lens of GP. The goal of the course is to equip students with the necessary knowledge and tools to apply probabilistic machine learning in their own work, be it theoretical or applied. Furthermore, students sharpen their general understanding of probabilistic and algorithmic reasoning.

- *Current*: While the course starts from well established foundations of exponential family distributions and GP regression, it progresses towards topics of current research interest, such as uncertainty in deep learning and efficient implementation of GPs.
- *Practical*: The course emphasizes algorithmic challenges posed by probabilistic machine learning and exercises guide students through efficient implementation of the relevant techniques.
- *Actionable*: The open structure of in-class discussion encourages students to explore the application of the course material to their own area of research as well and domain specific questions outside of machine learning research.

Evaluation

The course evaluation is based on two components:

1. Completion of weekly exercises: Students present their solutions to weekly exercises in class. These solutions should show an honest effort and active engagement with the material.

2. Active participation in the weekly meetings: This includes active participation in discussions through posing relevant questions, attempts to answer other students questions and connecting the material to real world problems, e.g. from their own field of research.

Materials

This course uses pre-recorded lectures from the courses *Probabilistic Machine Learning* and *Numerics of Machine Learning* that were offered by Philipp Henning in the summer term 2023 and winter term 2022 respectively at the University of Tübingen as part of a graduate program in Machine Learning. The recordings are available here and slides here. These materials are available under a CC BY-NC-SA 4.0 license. Selected chapters from the three volume series Probabilistic Machine Learning by Kevin Murphy as well as Phillip Hennig's textbook Probabilistic Numerics: Computation as Machine Learning will serve as supplementary reading.

Instructor

Brennan Klein is an associate research scientist at the Network Science Institute, with a joint affiliation at the Institute for Experiential AI. He is the director of the Complexity & Society Lab. His research spans two broad topics: 1) Information, emergence, and inference in complex systems — developing tools and theory for characterizing dynamics, structure, and scale in networks, and 2) Public health and public safety — creating and analyzing large scale datasets that reveal inequalities in the United States, from epidemics to mass incarceration. Dr. Klein received a PhD in Network Science in 2020 from Northeastern University and got his BA in Cognitive Science & Psychology from Swarthmore College in 2014. Website: brennanklein.com.

Schedule

The following schedule is tentative and might see revisions as the course progresses.

Week 1: Thur. Sep. 7, 2023 – Introduction & Reasoning under Uncertainty

Lectures:

- Lecture 1: Introduction
- Lecture 2: Reasoning Under Uncertainty

Supplementary Reading

- Murphy 2: 2.1 Probability: Introduction
- Murphy 2: 3.2 Bayesian Statistics
- Murphy 2: 4.2 Directed Graphical Models

Week 2: Thur. Sep. 14, 2023 – Continuous Variables & Exponential Families I

Lectures

- Lecture 3: Continuous Variables
- Lecture 4: Exponential Families

Supplementary Reading

- Murphy 2: 2.4 The Exponential Family
- Murphy 2: 2.5 Transformations of Random Variables
- Murphy 2: 3.4 Conjugate Priors

Week 3: Thur. Sep. 21, 2023 – Exponential Families II & Gaussian Probability Distribution

Lectures

- Lecture 5: Exponential Families II
- Lecture 6: Gaussian Probability Distribution

Supplementary Reading

- Murphy 1: 2.6 Univariate Gaussian (normal) Distribution
- Murphy 1: 3.2 The Multivariate Gaussian (normal) Distribution - 3.3 Linear Gaussian Systems
- Murphy 2: 2.3 Gaussian Joint Distribution

Week 4: Thur. Sep. 28, 2023 – Parametric Regression & Gaussian Processes (GP)

Lectures

- Lecture 7: Parametric Regression
- Lecture 8: Gaussian Processes

Supplementary Reading

- Murphy 1: 12 Generalized Linear Models
- Murphy 2: 15.1 GLM: Introduction - 15.2 GLM: Linear Regression
- Murphy 1: 17.2 Gaussian Processes

Week 5: Thur. Oct. 5, 2023 – Understanding through an Extensive Example

Lectures

- Lecture 9: Understanding Gaussian Processes
- Lecture 10: Gaussian Processes Regression: An Extensive Example

Supplementary Reading

- Murphy 2: 18.1 GP: Introduction - 18.5 GP with non-Gaussian Likelihoods

---

Week 6: Thur. Oct. 12, 2023 – Understanding GP through Kernels and Linear Algebra

Lectures

- Lecture 11: Understanding Kernels and Gaussian Processes
- Lecture 12: The Role of Linear Algebra in Gaussian Processes

Supplementary Reading

- Schölkopf & Smola Learning with Kernels (2002) Cahpter 1 A Tutorial Introduction
- Murphy 1: 7.6 Other Matrix Decompositions

Week 7: Thur. Oct. 19, 2023 – Computation, Inference & Logistic Regression

Lectures

- Lecture 13: Computation and Inference
- Lecture 14: Logistic Regression

Supplementary Reading

- Henning 2022 Chapter III.14 - III.20 Linear Algebra
- Murphy 1: 10 Logistic Regression
- Murphy 2: 12 Generalized Linear Models

Week 8: Thur. Oct. 26, 2023 – GP Regression & Deep Learning

Lectures

- Lecture 15: Gaussian Process Regression
- Lecture 16: Deep Learning

Supplementary Reading

- Murphy 1: 13 Neural Networks for Tabular Data
- Murphy 2: 16 Deep Neural Networks

Week 9: Thur. Nov. 2, 2023 – Probabilistic & Uncertain Deep Learning

Lectures

- Lecture 17: Probabilistic Deep Learning
- Lecture 18: Uncertainty in Deep Learning

Supplementary Reading

- Murphy 2: 17 Bayesian Neural Networks
- Murphy 2: 18.7 GPs and DNNs

Week 10: Thur. Nov. 9, 2023 – Use Cases & Gauss-Markov Models

Lectures

- Lecture 19: Uses of Uncertainty for Deep Learning
- Lecture 20: Gauss-Markov Models

Supplementary Reading

- Henning: I.5 Gauss-Markov Processes: Filtering and SDEs

Week 11: Thur. Nov. 16, 2023 – Parameter Inference

Lectures

- Lecture 21: Parameter Inference I
- Lecture 22: Parameter Inference II

Supplementary Reading

- Murphy 1: 8.7 Bound Optimization
- Murphy 2: 6.5 Bound Optimization
- Murphy 2: 10.1 Variational Inference

Week 12: Thur. Nov. 23, 2023 – Thanksgiving

Thanksgiving: No classes.

Week 13: Thur. Nov. 30, 2023 – Variational Inference & Historic Perspective

Lectures

- Lecture 23: Variational Inference
- Lecture 24: Historical Perspective

Supplementary Reading

- Murphy 2: 10 Variational Inference

Week 14: Thur. Dec. 7, 2023 – Probabilistic Numerics

Lectures

- Numerics of ML 6: Solving Ordinary Differential Equations
- Numerics of ML 7: Probabilistic Numerical ODE Solvers

Supplementary Reading

- Henning: VI Solving Ordinary Differential Equations

Week 15: Thur. Dec. 14, 2023 – Review Week

Recapitulate the most important concepts introduced during the course. Discuss the application of the course material to current and future research projects, as well as their implications for network science in general.
