M2794.0057 Advanced Topics in Dynamics, Control, and Robotics:
Machine Learning and Stochastic Estimation in Robotics
Fall 2014 Syllabus

Introduction: Methods from machine learning are starting to play a greater role in robotics, by enabling robots to learn from experience and to acquire skills needed to operate autonomously in unstructured environments. Some of these skills include sensorimotor capabilities such as locomotion, grasping, object recognition, manipulation, and the learning of complex tasks. The aims of this course are to learn the fundamentals of machine learning in a robotics context, and to examine the state-of-the-art in how machine learning algorithms are being used to enhance the operational capabilities of robots. Since machine learning has strong connections to probability and statistics and optimization, the early portion of the course will be devoted to the study of optimal estimation and stochastic systems.

Prerequisites: This course is a graduate advanced topics course, and students are expected to have previous exposure to the fundamentals of systems and control theory, and also to basic concepts in robot mechanics and control. An introductory graduate course in systems and control or equivalent, and an undergraduate course in robotics or equivalent, are necessary prerequisites. Students should also have studied probability and statistics as covered in an undergraduate applied mathematics course. The course also draws heavily upon methods of linear algebra, and as such courses on multidimensional systems analysis and optimization, while not required, will be helpful. **All students should obtain prior permission of the instructor at the beginning of the course**.

Course Instructors: The instructors for this course are Frank Chongwoo Park (fcp@snu.ac.kr) and Yung-Kyun Noh (yungkyun.noh@gmail.com); instructor profiles are provided at the end of this document. F.C. Park's office is located in Building 301; Y.-K. Noh's office is located in Building 302, Room 413. Meetings can be arranged by previous appointment.

Lectures: The course format will primarily be based on two weekly lectures given by the instructors. The lectures will be given on Tuesdays and Thursdays from 14:00–15:15 in Building 301, Room 204. F.C. Park will give lectures on optimal estimation and stochastic systems, while Y.-K. Noh will give lectures on the foundations of machine learning. **All the lectures will be given in English.** Towards the end of the course, select guest lectures covering advanced topics and applications may also be offered.

Another important component of the course will be videotaped student lectures given by the participants. One of the best ways for a student to learn new material is to teach it to others, and as such, each student will be asked to present a 20-minute lecture on a basic subtopic of machine learning. The lectures will be videotaped in our multimedia studio, and made available for online viewing and critique by registered students of the course. Topics will be assigned and the lectures videotaped during the second half of the course.

Course Webpage: A course webpage will be maintained at *http://etl.snu.ac.kr*. All lecture notes, homework assignments, solutions, and announcements will be made available on the course webpage.

Course Materials: This course will draw upon a number of educational resources and materials. The material on optimal estimation and stochastic systems will primarily be based on a set of unpublished lecture notes by Roger Brockett on stochastic control. For the material on machine learning, the lectures will be based on selected chapters from "Machine Learning: A Probabilistic Perspective" by Kevin P. Murphy (MIT Press) and "Pattern Recognition and Machine Learning" by Christopher M. Bishop (Springer). For reinforcement learning, "Probabilistic Robotics" by Sebastian Thrun et al. (MIT Press) will be the primary reference. Supplemental lecture notes, papers, and handouts will be distributed to the class as necessary throughout the course.

Grading: The grading for the course will be based on approximately 4 problem sets (20%), a midterm exam (20%), a final exam (25%), evaluation of the student online lecture (10%), a design project (20%), and class participation (5%). Note that the weights are approximate, and may be adjusted accordingly at the discretion of the instructors. Details of the design project will be given later in the course, and will involve an implementation and analysis of a machine learning algorithm to a current problem of topical interest in robotics or a related area.

Tentative Sequence of Topics: We will try to adhere to the following sequence of topics (topics and schedule may be adjusted depending on circumstances):

1. **September 2: Introduction and probability review**
   - Course introduction
   - Probability review
2. **September 4: Linear dynamic systems**
   - Linear dynamic systems
   - Controllability and observability
   - Covariance matrix and the propagation of errors
3. **September 9: HOLIDAY (NO LECTURE)**
4. **September 11: Linear filtering**
   - Optimal linear filtering
   - Kalman filter in continuous and discrete time
5. **September 16: IROS CONFERENCE (NO LECTURE)**
6. **September 18: IROS CONFERENCE (NO LECTURE)**
7. **September 23: Nonlinear estimation**
   - Nonlinear minimum variance estimation
   - Nonlinear least-squares estimation
8. **September 25: Jump process stochastic differential equations**
   - Poisson counters and finite-state continuous time jump processes
   - Ito calculus for Poisson counters

9. **September 30: Wiener process stochastic differential equations**
   - Brownian motion and stochastic differential equations
   - Ito calculus for Wiener processes
   - Fokker-Planck equations for the density
10. **October 2: Nonlinear filtering**
    - Hidden Markov models
    - Conditional density equations
11. **October 7: Machine learning introduction**
    - Issues in learning from data
    - Generative approach for machine learning
12. **October 9 HOLIDAY (NO LECTURE)**
13. **October 14: Parametric and nonparametric methods**
    - Density estimation : parametric and nonparametric methods
    - Gaussian density models
14. **October 16: Parameter estimation for Gaussian models**
    - Maximum likelihood and maximum a posteriori parameter estimation for Gaussian models
15. **October 21: Kernel estimation**
    - Kernel density estimation and nearest neighbor methods
    - Issues for multivariate data
16. **October 23: MIDTERM EXAMINATION**
17. **October 28: Graphical models**
    - Graphical models
18. **October 30: Graphical inference**
    - Conditional probability and marginal probability
    - Basic inference methods in graphical models
19. **November 4: Supervised learning**
    - Supervised learning
    - Generative methods for classification and regression
20. **November 6: Reinforcement learning**
    - Reinforcement learning
    - Markov decision processes

21. **November 11: Value function approximation**
    - Value function approximation
22. **November 13: Gaussian processes**
    - Gaussian model revisited
    - Ideas of Gaussian processes
23. **November 18: Gaussian process regression**
    - Gaussian processes
    - Gaussian process regression
24. **November 20: Nonlinear kernel models**
    - Nonlinear extension of linear models using kernels
    - Function-space view and weight-space view
25. **November 25: Embeddings, Gaussian process models**
    - Nonlinear embedding of data within manifold
    - Locally linear embedding and Isomap
    - Gaussian process latent variable model
    - Gaussian process dynamical model
26. **November 27: Harmonic maps and learning**
    - Harmonic maps and embedding models
27. **December 2: Control with Gaussian processes**
    - Control with Gaussian processes
28. **December 4: Deep learning**
    - Guest lecture on deep learning
29. **December 9: Course review**
    - Course review
30. **December 11: FINAL EXAMINATION**

Course Instructor Profiles: **F.C. Park** is a professor of mechanical engineering at Seoul National University. He received his B.S. in electrical engineering from MIT in 1985, and Ph.D. in applied mathematics from Harvard University in 1991. His research interests are in robot mechanics, planning and control, mathematical systems theory, computer vision, and related areas of applied mathematics. Further details are available at http://robotics.snu.ac.kr.

Yung-Kyun Noh is a research professor in the department of Computer Science at KAIST. He received his B.S. in Physics from POSTECH in 1998, and his Ph.D. in Computer Science from the Interdisciplinary Program in Cognitive Science at Seoul National University in 2011. His research interests are in metric learning and dimensionality reduction in machine learning. He is especially interested in applying statistical theory of nearest neighbors to real and large datasets. Further details are available at http://ailab.kaist.ac.kr.