



---

## SF2957 HT19 Statistical Machine Learning 7.5 credits

**Lecturer:** Henrik Hult. Email: hult@kth.se, phone: 08-790 6911

**Course literature:** Lecture notes and extracts from various books

- Casella, G. and Berger, R.L., *Statistical inference*, 2nd Ed., Duxbury Thomson Publishing, 2002.
- Goodfellow, I., Bengio, Y., and Courville, A., *Deep learning*, MIT Press, 2016. Available electronically at:
- Mackay, D., *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003. Available as PDF at:
- Rasmussen, C.E. and Williams, K.I., *Gaussian processes for machine learning*, MIT Press, 2006. Available as PDF at:
- Sutton, R.S. and Barto, A.G., *Reinforcement learning: An introduction*, MIT Press, 2017. Available as PDF at:
- Kushner, H. J. and Yin, G. G., *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd Ed., Springer, 2003.

**Learning outcomes:** This course presents an overview of advanced methods of statistical machine learning. The course addresses both theoretical and practical aspects of statistical machine learning. Computer-aided projects with a variety of datasets forms an essential learning activity. To complete the course, the student must be able to:

- formulate and apply statistical decision theory
- formulate and apply advanced methods in statistical machine learning
- design and implement advanced methods in statistical machine learning for applications

**Examination:** The examination consists of projects (3.0 credits) and a final exam (4.5 credits). Grades are given in the range A-F, and Fx. The grade Fx gives you the right to a complementary examination to reach the grade E. The criteria for Fx is F and that an isolated part of the course can be identified where you have shown particular lack of knowledge and that the examination after a complementary examination on this part can give the grade E.

**Grading:** The requirements for higher grades are as follows.

Learning outcome	E	C	A
<b>Formulate and apply statistical decision theory</b>	Formulate concepts in statistical decision theory and apply the concepts to solve problems	Motivate methods in statistical decision theory and their properties	Generalise concepts in statistical decision theory and derive fundamental relations
<b>Formulate and apply advanced methods in statistical machine learning</b>	Formulate advanced methods in statistical machine learning and apply the methods to solve problems	Motivate advanced methods in statistical machine learning and their properties	Generalise advanced methods in statistical machine learning and derive fundamental relations
<b>Design and implement advanced methods in statistical machine learning for applications</b>	Ability to design advanced methods in statistical machine learning, implement the methods by writing computer programs to solve applied problems and present, in writing, the methods, motivation of the methods, and results. Applies to grade E only.		

**Projects:** The projects examine the third learning outcome. Project information will be available on the course's Canvas page. Participation in a project group and writing a report is mandatory to pass the course.

- You must work in groups of at most four people.
- You must work actively on the project together with the other group members and be prepared to account for your own and other group members' contribution.
- You must hand in the report by uploading it on canvas before the deadline.
- You must, together with your project group, be prepared to make a presentation of your project or part of your project.

**Final Exam:** The final exam takes place on Jan 7, 2020, 08:00-13:00. For the final exam you may bring one page (size A4) of written/typed notes on both sides (cheat-sheet).

The final exam consists of three parts.

**Part I** consists of two problems and examines the first learning outcome, statistical decision theory. Each of the problems give at most 10 points. has a theory part (a) that gives 4 points and a problem part that gives 6 points. You need 10 points to pass part I.

**Part II** consists of two problems and examines the second learning outcome, statistical machine learning. Each of the problems give at most 10 points. has a theory part (a) that gives 4 points and a problem part that gives 6 points. You need 10 points to pass part II.

To pass the final exam you must pass part I and part II.

**Part III** consists of one problem that examines higher level learning outcomes, typically constructed to combine concepts or generalise concepts from the course.

**Grade limits:** To pass the course and receive grade E, it is required to pass the projects and part I and part II on the final exam. Given that a passing grade is obtained, higher grades will be based on the score of the final exam according to the following:

A: 41-50

B: 35-41

C: 29-34

D: 22-28

E: 20-21

FX: 8-9 points on part I and 10+ points on part II or 10+points on part I and 8-9 points on part II