Aalto University
School of Science

CS-E4070 — Computational learning theory

Slide set 00 : course logistics and syllabus

Cigdem Aslay and Aris Gionis

Aalto University

spring 2019

course logistics

- lectures : mon 2-4pm and wed 10-12
  - wednesdays are normal lectures, no exercises
  - no lectures on mon 23/4 and wed 1/5
- room : T6 at T-building (except today)
- instructors :
  - Cigdem Aslay, cigdem.aslay@aalto.fi
  - Aris Gionis, aristides.gionis@aalto.fi
- office hours :
  - by arrangement (email)
- other communication :
  - mycourses discussion forum or email

textbooks and reading material

- Michael Kearns and Umesh Vazirani (K&V)
  An Introduction to Computational Learning Theory
  The MIT Press, 1994
  Available as an e-book via the Aalto online library

- Shai Shalev-Shwartz and Shai Ben-David (SS&BD)
  Understanding Machine Learning: From Theory to Algorithms
  Cambridge University Press, 2014
  Available online

- "Classics" papers listed in mycourses

workload

passing the course and the final grade will depend on

1. three (3) homeworks
   - each homework will be a 'pen and paper' type problem set
   - return via mycourses, typed in Latex
2. studying a recent research paper
   - select a paper from COLT conference, 2010 onwards
   - propose your paper by May 1 via a short document in mycourses, motivating your selection
   - study in detail, you may need to read other key papers
   - write a paper summary and return via mycourses by June 1

- homeworks and paper summary should be done individually (no group work)

motivation for the course

- machine learning is a thriving topic
- a wealth of forefront methods
- important applications in the spotlight
  - from machine translation to recommendation systems to self-driving cars
- beyond novel methods and applications we are interested in fundamental computational questions
  - what can be learned by machines?
  - what error bounds can be achieved?
  - what computational resources are required?
  - what cannot be learned?

why study computational learning theory?

- obtain rigorous treatment of fundamental questions
- study foundations of machine learning
- motivate practical machine-learning methods
- seek for understanding of successful heuristics
- enrich your machine-learning curriculum

disclaimers

1. emphasis on theory

   no programming

   will not learn about the latest models that will help you find a data-science job, or build your own start up

2. your lecturers are not experts on the topic

   you may want to think of the course as a study group

list of topics and tentative schedule

lec 1	basic notions and introduction to PAC learning
lec 2	uniform and non-uniform learnability, Occam's razor
lec 3	VC dimension and sample complexity
lec 4	Rademacher complexity and covering numbers
lec 5	weak and strong learning
lec 6	learning in the presence of noise
lec 7	submodular optimization
lec 8	applications of submodularity in machine learning
lec 9	online learning: mistake-bound models
lec 10	online learning: no-regret models

prerequisites

- basic discrete math and probability
- big-$\mathcal{O}$ notation and basic analysis of algorithms
- familiarity with mathematical proof principles
- basic knowledge of the theory of **NP**-completeness