

# CSE 547: Machine Learning for Big Data

**Instructor:** Tim Althoff

**Lectures:** Tuesday/Thursday 10:00-11:20am PDT in CSE2 G10

**Course website:** <https://courses.cs.washington.edu/courses/cse547/23wi/>

**Contact:**

- Use Ed to post questions: <https://edstem.org/us/courses/32320/discussion/>.
- For emergencies and personal questions, e-mail us at [cse547-instructors@cs.washington.edu](mailto:cse547-instructors@cs.washington.edu)

**TAs and office hours:** See the course website for times and locations.

## Topics

- MapReduce and Spark/Hadoop
- Frequent itemsets and Association rules
- Near Neighbor Search in High Dimensions
- Locality Sensitive Hashing (LSH)
- Dimensionality reduction: SVD and CUR
- Recommender Systems
- Clustering
- Analysis of massive graphs
- Link Analysis: PageRank, HITS
- Web spam and TrustRank
- Proximity search on graphs
- Large-scale supervised machine learning
- Mining data streams
- Optimizing submodular functions

## Assignments and grading

- 10 Colabs (20%)
- Four problem sets requiring coding and theory (40%)
- Final project (40%)
- Extra credit: Ed and course participation, reporting bugs in course materials (up to 2%)

## Homework policy

**Questions** We try very hard to make questions unambiguous, but some ambiguities may remain. Ask (i.e., post a question on Ed) if confused, or state your assumptions explicitly. Reasonable assumptions will be accepted in case of ambiguous questions.

**Honor code** We take honor code extremely seriously:

(<https://www.cs.washington.edu/academics/misconduct>).

We strongly encourage students to form study groups. Students may discuss and work on homework problems in groups. However, each student must write down the solutions and the code independently. In addition, each student should write down the set of people whom they interacted with.

**Late assignments** Each student will have a total of two late periods to use for homeworks. A late period lasts 48 hours from the original homework deadline. No assignment will be accepted more than one late period after its due date.

**Assignment submission** All students must submit their homeworks via Gradescope (<http://www.gradescope.com>). Students can typeset or scan their homeworks. Scanned homeworks must be absolutely clear and legible. Points may be deducted for unclear parts of handwritten homework.

Students also need to upload their code on Gradescope. Put all the code for a single question into a single file and upload it. Refer to the course FAQ for more info.

**Regrade requests** We take great care to ensure that grading is fair and consistent. Since we will always use the same grading procedure, any grades you receive are unlikely to change significantly. However, if you feel that your work deserves a regrade, submit your request within a week of receiving your grade on Gradescope. However, note that we reserve the right to regrade the entire assignment.

## Prerequisites

Students are expected to have the following background (recitation sessions will refresh these topics):

- The ability to write non-trivial computer programs (at a minimum, at the level of CSE 332, CSE 373, or equivalent). Good knowledge of Python/Java will be extremely helpful since most assignments will require the use of Hadoop/Java.
- Familiarity with basic probability theory is essential (any introductory probability course).
- Familiarity with writing rigorous proofs (e.g. CSE 311 or equivalent).
- Familiarity with basic linear algebra (e.g. MATH 308).
- Familiarity with algorithmic analysis (e.g. CSE 417, CSE 421).

## Materials

Notes and reading assignments will be posted on the course web site. Readings for the class will be from:

- Mining Massive Datasets by J. Leskovec, A. Rajaraman, J. Ullman (PDFs at <http://mmds.org>).

## Important dates

Assignment	Out Date	Due Date (all 23:59pm)
Assignment 1	Jan 5	Jan 19
Assignment 2	Jan 19	Feb 2
Project proposal		Jan 26
Assignment 3	Feb 2	Feb 16
Project milestone		Feb 12
Assignment 4	Feb 16	Mar 2
Final report		Mar 12
Final poster		Jun 13

We will also hold three review sessions in the first two weeks of the course:

- Spark tutorial and help session. Thursday, Jan 5, 3:30-5:00pm in CSE2 371.
- Review of basic probability and proof techniques. Tuesday, Jan 10, 3:30-5:00pm in CSE2 371 201.
- Review of basic linear algebra. Thursday, Jan 12, 3:30-5:00pm in CSE2 371.

## Next steps for students

- Register for Ed: <https://edstem.org/us/courses/32320/discussion/>.
- Register for Gradescope: <https://www.gradescope.com/courses/480663> course code 7G5RY2
- Download Spark (VM), start Homework 0