

CSC 466: Knowledge Discovery from Data (KDD)  
Fall 2021  
Course Syllabus

September 12, 2021

**Instructor:** Alexander Dekhtyar  
**email:** dekhtyar@calpoly.edu  
**office:** 14-212 (new office location)

What	Who	When	Where	
Lecture	Section 01	TR	9:40 am – 11:00am	034-228 (Dexter)
	Section 03	TR	1:40 pm – 3:00pm	052-0E27 (Spider)
Lab	Section 01	TR	12:10pm – 1:30pm	14-301
	Section 03	TR	1:40pm – 3:00pm	14-301
<b>Final Exam</b>	Section 01	December 9 (Thursday)	10:10 – 1:00pm	034-228
	Section 03	December 7 (Tuesday)	1:10 – 4:00pm	052-0E27

*Note: the class will not have a written final exam, but we will most likely use the exam time for team project presentations.*

**Office Hours**

	When	Where	Zoom
Monday	9:10am - 11:00am	14-212	<a href="https://calpoly.zoom.us/my/dekhtyar">https://calpoly.zoom.us/my/dekhtyar</a>
Tuesday	5:10pm - 6:00pm	14-212	<a href="https://calpoly.zoom.us/my/dekhtyar">https://calpoly.zoom.us/my/dekhtyar</a>
Wednesday	9:10am - 10:00am	14-212	<a href="https://calpoly.zoom.us/my/dekhtyar">https://calpoly.zoom.us/my/dekhtyar</a>

Additional appointments can be scheduled by emailing the instructor at [dekhtyar@calpoly.edu](mailto:dekhtyar@calpoly.edu).

## Description

This class is an overview of the field of knowledge discovery from data (KDD, also often referred to as "Data Mining" or "Machine Learning") and related technologies. The course is intended for junior/senior students in Computer Science, Software Engineering and Computer Engineering majors, as well as for students completing the Cross-Disciplinary Studies Minor in Data Science. The

course gives a broad overview of **data mining** (association rules mining, classification, clustering), **information filtering** and recommender systems, **information retrieval** and web search, and **web mining**.

## Learning Objectives

After taking the course the students are expected to be able to

1. **recognize** different types of KDD procedures and **identify** their uses;
2. **implement** algorithms/methods/techniques for KDD tasks to **solve** KDD problems;
3. **interpret** and **analyze** the *results* of KDD processes;
4. **recognize** and **evaluate** *societal impact* of KDD technology, **make informed choices** about use of KDD technology.

## Textbook

- Bing Liu. *Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data*, Springer, 2nd ed. 2011. ISBN: 978-3642194597.

This book is **mandatory**. It contains almost all material studied in the course (and then some). While we rely on instructor's lecture notes as much as we rely on the content of the book, the book is extremely useful.

In addition, some other books may be of use. If you want an alternative take on most of the material covered in class, I recommend (as an optional book)

- Mohammed J. Zaki, Wagner Meira Jr. *Data Mining and Analysis: Fundamental Concepts and Algorithms*, Cambridge University Press, 2014, ISBN: 978-0-521-76633-3.

This book is the textbook in Professor Khosmood's version of CSC 466.

## Topics

No.	Topic	Duration	Liu	Zaki,Meira
1.	Association Rules	1	Chapter 2	Chapters 8, 9
2.	Supervised Learning (Classification)	2	Chapter 3	Chapters 18, 19, 22
3.	Unsupervised Learning (Clustering)	2	Chapter 4	Chapters 13, 14, 15, 17
4.	Collaborative Filtering	1	Chapter 11, 12	
5.	Information Retrieval	2	Chapter 6	
6.	Link Analysis	2	Chapter 7	

Please note that the order in which these topics are covered may be different than the order in which they are presented above.

## Grading

<b>Labs and Homeworks</b>	60%
<b>Lab Exam</b>	10%
<b>Projects</b>	30%

## Course Policies

### Prerequisites

The official prerequisite for this course is **CSC 349 (Algorithms)**. This prerequisite is enforced **strictly**. CSC 466 can be viewed as an advanced algorithms course for a certain important category of algorithms. Therefore, it is important that the craft of algorithm design is not a mystery to anyone in the class.

### Exams

The course will have **no written exams**. In their stead you will be offered three things:

1. A **take-home group project**. The project assignment will come out after Week 5-6 of the class, once the finish covering the core topics of the course (classification and clustering). The group assignment will ask you to apply the methods studied in the course to exploratory analysis of some real-life data.
2. A **take-home individual assignment**. This assignment will come out in the last 2-3 weeks of the class.
3. An **in-class lab exam**. The lab exam will most likely take place during the lab periods of Week 10 of the course. You will be asked to write some code to implement some straightforward KDD tasks. We will be using Jupyter Notebooks for the exam, and you will be given some sample problems at least a week or two ahead of the time.

Most of lab assignments and other coursework in this course is done in either small teams or in pairs. The individual assignment and the lab exam were introduced into the class to allow me to properly assess individual abilities of students, that otherwise may have been obscured due to teamwork.

The reserved final examination time will be used for group project presentations.

### Labs, Homeworks

Hands-on KDD-related activities are the core part of the course. Some activities will be set up as lab exercises, some other activities may be offered as purely take-home assignments (this will be determined by the pace of the course).

Each lab/homework assignment will involve some data analysis task, that may involve using existing software, software provided by the instructor as well as (and mostly) the software developed by you. The course concentrates on **basic algorithms for performing standard KDD tasks**: the labs/assignments give you an opportunity to cement the knowledge of the algorithms covered in class.

Most of the lab assignments are pair programming assignments, although some exceptions may be made either in favor of small teams, or in favor of individual assignments. This will be announced ahead of each lab assignment.

**Note:** Machine learning, data mining and other KDD algorithms that we are going to study in the class are widely available, both as open source code and, in some cases, as methods/functions in popular KDD libraries. **One of the goals of this course is to have you implement these algorithms from scratch!** The assignments will specify when you can, and when you cannot third-party code/libraries to achieve the goals of the assignment.

**Note 2:** At the same time, in a lot of the assignments, the main deliverable will be not the code you write, but rather, the insight you obtain by running your code on the data provided to you. Please be aware of that, as this shift in what is the main deliverable, is perhaps one of the key unique features of CSC 466.

### Late Submissions

*Late lab and assignment submissions* are strongly discouraged. The course will run on a tight schedule, and not submitting on time will lead to time carved out of the next assignment. A penalty of 10 - 30% will be assessed for any submissions that are late by less than 24 hours. No credit will be given for any later submissions. You are encouraged to submit on time even if your submission it is not perfect. You can then resubmit a fixed version late, subject to the abovementioned rules. When more than one submission is present, we will independently grade two submissions: (i) the latest on-time submission and (ii) the latest late submission for which non-zero credit can be assessed. Your grade for the project will be the **maximum** of the two grades.

### Course Calendar

Please note the following changes to our schedule:

- **September 21, Tuesday** (first day of the class): double lecture. No lab.
- **September 23, Thursday**: no class. Feel free to use the lab during the lab period, but the lab period is unsupervised.
- **November 11, Thursday**: Veterans Day. No class.

I am attending a virtual conference during the first week of classes, this necessitates the change in schedule for Thursday, September 23, and a double lecture (to compensate) on Tuesday, September 21.

## Communication

The class will have an official mailing list. The email address for the mailing list is *csc-466-01-2198@calpoly.edu*. All students enrolled in the class are automatically subscribed to the mailing list (using the email address that the CS department has on file).

I encourage questions during classtime and questions via email. My answers to email questions may be broadcast to the entire class via the mailing list, if the answer may be relevant to everyone (e.g. a correction in a text of a handout, or a clarification of a homework problem), and may also appear on the web page. The questions can also be posted to the mailing list directly. The mailing list will also be used for all announcements related to the course. It is your responsibility to read your class-related email. Failure to read email posted to the mailing list cannot be used as an excuse in the class.

## Office Hours

While we are under COVID-19 restrictions, all office hours will be held in hybrid mode. That is, office hours will be held in person in my office, 14-212, with Zoom turned on, in case someone wants to talk to me without having to sit in my office. The Zoom address is

<https://calpoly.zoom.us/my/dekhtyar>

or

<https://calpoly.zoom.us/j/2118158081>

It will also be sent to you via email.

On some rare occasions, office hours may be fully virtual over zoom. These occasions will be announced ahead of time for your convenience.

## Web Page

Class web page can be found at

<http://www.csc.calpoly.edu/~dekhtyar/466-Fall2021>

Through this page you will be able to access all class handouts including homeworks, lab assignments, project information, lab/project data and lecture notes.

Links to additional information, and notes and announcements will also be posted.

Please note, this class does not use Canvas.

## Slack

In addition to the course mailing list and course web page, we will also have Slack workspace. You should have received invitations sent to your Cal Poly

email. I have Slack open on all my computing devices, which means that if you need to discuss something with me, this may be a better medium than email, whether this is a one-on-one conversation, or a conversation that involves and entire section or a group of students.

## **COVID-19 Policies**

This course will follow the University's COVID-19 policies, regulations, and recommendations.

Specifically:

- Your physical presence in class is a self-certification that you are following the procedures proscribed by the University w.r.t. COVID-19 testing and self-assessment.
- You **must** wear a mask to the class. I will try to have a supply of disposable masks available in case you did not bring one with you, and I will offer you one if I see that you are not wearing a mask.
- If you refuse to wear a mask, I will ask you to leave the class. If you refuse to leave the class, per University recommendations, I may decide to end the class early/cancel the class for the day. If this behavior repeats, it will be reported to the Office of Student Affairs.
- I will be lecturing while wearing a mask. If my voice is not sufficiently loud, please let me know, I will raise it. If what I am saying is muffled, please let me know, I will make every effort to speak more clearly (subject to my natural accent).
- If you have contracted COVID-19 and are isolating/in quarantine, or if you are unable to attend classes for any other reason, please let me know. I am available via zoom during all office hours (see above), and if necessary, we may schedule a zoom conversation outside of office hours to discuss coursework.
- If you are absent for more than two weeks due to COVID-19 or any other unforeseen circumstances, I will work with you in case you want to withdraw from the course. Absences of more than two weeks from the course may create a situation where you are in danger of falling behind in the coursework and not recovering. If you are interested in retaking CSC 466, I will make sure you will be registered for another section of it in Winter or Spring quarters (although it might be with a different instructor).

## **Academic Integrity**

### **University Policies**

Cal Poly's Academic Integrity policies are found at

<http://www.academicprograms.calpoly.edu/academicpolicies/Cheating.htm>

In particular, these policies define *cheating* as (684.1)

*“... obtaining or attempting to obtain, or aiding another to obtain credit for work, or any improvement in evaluation of performance, by any dishonest or deceptive means. Cheating includes, but is not limited to: lying; copying from another’s test or examination; discussion of answers or questions on an examination or test, unless such discussion is specifically authorized by the instructor; taking or receiving copies of an exam without the permission of the instructor; using or displaying notes, ”cheat sheets,” or other information devices inappropriate to the prescribed test conditions; allowing someone other than the officially enrolled student to represent same.”*

Plagiarism, per University policies is defined as (684.3)

*“... the act of using the ideas or work of another person or persons as if they were one’s own without giving proper credit to the source. Such an act is not plagiarism if it is ascertained that the ideas were arrived through independent reasoning or logic or where the thought or idea is common knowledge. Acknowledgement of an original author or source must be made through appropriate references; i.e., quotation marks, footnotes, or commentary.”*

University policies state (684.2): “Cheating requires an “F” course grade and further attendance in the course is prohibited.” (appeal process is also outlined, see the web site above for details.). Plagiarism, per university policies (684.4) can be treated as a form of cheating, although a level of discretion is given to the instructor, allowing the instructor to determine the causes of plagiarism and effect other means of remedy. It is the obligation of the instructor to inform the student that a penalty is being assessed in such cases.

### **Course Policies**

All homeworks are to be completed by each student **individually**. Lab assignments are to be completed by the appropriate units (individual, pair, group), and no code/solution-sharing between units is permitted. Students are encouraged to discuss class content among themselves but NOT in a manner that constitutes plagiarism and cheating as defined above (e.g., you can solve together a problem from the textbook that had not been assigned in the homework, but you should solve assigned problems individually).