



California Institute of Technology  
Department of Computing + Mathematical Sciences

**IDS/ACM/CS 157 Statistical Inference**

Spring 2026

**Lectures:** Tue & Thu 9:00am-10:25am in Baxter Lecture Hall (BLH)  
**Instructor:** Konstantin (Kostia) Zuev  
**Office:** Annenberg 114  
**Email:** [kostia@caltech.edu](mailto:kostia@caltech.edu) (please include “157” in the subject line)  
**Office Hour:** Tue 1pm-2pm, or by appointment (please, send an email to schedule)  
**Head TA:** Harsh Gandhi ([hgandhi@caltech.edu](mailto:hgandhi@caltech.edu))  
**TAs and OHs:** <https://piiazza.com/caltech/spring2026/idsacmcs157/staff>  
You are welcome to attend as many office hours as you like.

**Course Goals**

Statistical Inference is a branch of Mathematical Engineering that studies ways of extracting reliable information from limited data for learning, prediction, and decision making in the presence of uncertainty. The main goals of this course are:

- Develop statistical thinking and intuitive feel for the subject,
- Introduce the most fundamental ideas, concepts, and methods of Statistical Inference, and
- Explain how and why they work, and when they don’t.

If you do well in the class, you should be able to read (and understand) most contemporary papers that use statistical inference and perform statistical analysis yourself.

**Prerequisites**

This is an introductory course on statistical inference. No prior knowledge of statistics is assumed. However, a solid understanding of Probability is required. Ma 3 or ACM/EE/IDS 116 (or equivalent) is a “hard” prerequisite. A key part of the course is problem sets, where you will get experience in using the learned methods and models in applications via simulations in MATLAB. So, some familiarity with MATLAB (and programming in general) is desired, but this is a “soft” prerequisite: MATLAB is easy to pick up on the fly, especially for the purposes of this course.

**Textbooks**

The course is based on the following book, which is available in the SFL Library:

K.M. Zuev (2026), *Fundamentals of Statistical Inference: Foundations of Data Analysis*

The e-book will also be available on Piazza under “Resources”.

The hard copy is available on the Springer website: <https://tidd.ly/49yToBA>

**Grading**

Your final grade will be based on your Total Score. Your Total Score is a weighted average of Problem Sets (60%), Midterm Exam (20%), and Final Exam (20%). You can increase your Total Score by up to 5% if you participate actively in Piazza discussions in the Q&A section<sup>1</sup>. Every answer submitted before TAs or instructor answer, which is later endorsed as a “good answer” by TAs or instructor, gets 1% of the total score. There are no fixed thresholds for grades, but if your total score is 90% (80%, 70%, 60%), you are guaranteed at least “A” (“B”, “C”, “D”).

**Problem Sets**

There will be six Problem Sets. Problems (and solutions) will be posted on Piazza. For assignment and due dates see “Important Dates” below. Late submissions will not be accepted, but the Problem Set with the lowest score will be dropped and not counted toward your total score. Submitting wrong files or files in a wrong format is considered as a late submission. Extensions may be granted for academic, personal, or medical reasons. For extensions, please email the Head TA.

**Exams**

There will be two exams:

1. Midterm Exam: based on Lectures 1-8, **take-home**, 3h long, timed on Gradescope.
2. Final Exam: cumulative, based on all course material, **in-person**, 1h 15 min long, paper-based (no electronic devices). The Head TA will provide a review session before each exam. Both exams are closed-book but open-notes (*your* notes): only material written or typed by you may be used during exams. Electronic devices may be used only for typing and for arithmetic operations on the take-home midterm exam. The final exam is paper-based: no electronic devices are permitted.

---

<sup>1</sup> If you are interested in being a TA next year, try to be active on Piazza and help other students by answering their questions.

## Final Exam Policy

The final exam is an in-person exam and must be taken at the scheduled time. Students who are unable to take the exam due to illness, emergency, or other extenuating circumstances must contact the Deans' Office as soon as possible. Any exceptions or alternate arrangements require approval from the [Deans' Office](#) or [CASS](#). I am not able to accommodate individual rescheduling requests outside of these processes. Failure to take the final exam as scheduled, without prior approval from the Deans' Office, may result in a missing exam grade.

## Ethical Use of AI

You can use AI tools (e.g., ChatGPT) to support your learning in this course, but only in ethical and responsible ways. For example, it is fine to use AI to generate a practice exam based on the topics covered in the course. However, using AI to directly solve your problem sets or exams, to give you hints, or check your solutions for correctness is not allowed, as it undermines your learning and violates Caltech's Honor Code. When in doubt, ask yourself: would it be acceptable for a tutor to do this for you? If not, then it is also not appropriate to ask an AI to do it. Most importantly, keep in mind that you are here to train your own neural network, not the artificial one.

## Collaboration Policy

A detailed collaboration policy is given on the course website at:

<http://www.its.caltech.edu/~zuev/teaching/2026Spring/CollaborationIDS157.pdf>

In general, collaboration is encouraged everywhere except for the exams. Let's help each other and learn together!

If you get stuck with a homework problem, I encourage you to discuss it with other students (offline or online on Piazza).

But remember that you will have to prepare and submit your solution by yourself. No collaboration is allowed on the exams.

## Important Dates (All times are Pacific Times)

	Available	Due
Problem Set 1	1pm Tue, Apr 07	9pm Tue, Apr 14
Problem Set 2	1pm Tue, Apr 14	9pm Tue, Apr 21
Problem Set 3	1pm Tue, Apr 21	9pm Tue, Apr 28
Head TA Review	9am Tue, Apr 28	
Midterm Exam	1pm Tue, Apr 28	9pm Tue, May 05
Problem Set 4	1pm Tue, May 05	9pm Tue, May 12
Problem Set 5	1pm Tue, May 12	9pm Tue, May 19
Problem Set 6	1pm Tue, May 19	9pm Tue, May 26
Head TA Review	9am Tue, May 26	
Final Exam (in-person) Location: BLH / ANB 105 <sup>2</sup>	Start: 9am Thu, June 4	End: 10:15am Thu, June 4

## Websites

- Course website: <http://www.its.caltech.edu/~zuev/teaching/2026Spring/IDS157.htm>
- Problem sets, data sets, solutions, announcements, and class discussions will be managed via Piazza, which is designed such that you can get a quick help from your classmates, TA(s), and instructor. Instead of emailing questions to the teaching staff, I encourage you to post your questions on Piazza because
  - You will get the answers faster
  - Your classmates may also benefit from seeing the answers to your questions.

Here is the Piazza page: <http://www.piazza.com/caltech/spring2026/idsacmcs157/home>

- Problem sets and exams will be graded via Gradescope.
  - If you are a **registered student**, you will be enrolled on Gradescope by the end of the 1<sup>st</sup> week of classes, and you will receive a notification from Gradescope about your enrollment.
    - Please make sure that the email that you use on Gradescope is your official Caltech email.
  - If you are a **registered student** but have not been enrolled on Gradescope by the end of the 1<sup>st</sup> week of classes, please email the Head TA as soon as possible and ask to enroll you to Gradescope. Your absence on Gradescope means that, according to my records, you are not registered for the course.
  - If you want just to **audit the course**, it is fine, you will have access to Piazza and all course materials there (please email me and I will enroll you on Piazza), but you will not have access to Gradescope and your submissions will not be graded. If you audit the course this term, you should not register for the course in the future.

<sup>2</sup> Room ANB 105 is reserved for students with CASS accommodations. Baxter Lecture Hall is for students without CASS accommodations.

To submit your solution via Gradescope, you need to create a single PDF (not images) that contains the whole solution (for example, by scanning your solution), and then upload it to Gradescope. Here is a useful link:

- How can I submit my homework as a PDF?

<https://guides.gradescope.com/hc/en-us/articles/21862105254413-How-can-I-submit-my-homework-as-a-PDF>

Should you have any questions regarding Gradescope, please ask on Piazza: we will have many experts there.

### Suggested Study Process

To get the most out of IDS 157, here is my suggestion on the study process<sup>3</sup>:

- **Have Enough Sleep:** Good sleep is an important prerequisite for learning.
- **Attend Lectures:** Focus on understanding the big picture of what is going on.
- **Review the Relevant Book Sections:** Ideally on the same day as the lecture, and make sure everything is clear.
- **Ask and Answer Questions:** If something is not clear, ask on Piazza; help your classmates by answering their questions.
- **Summarize in Your Own Notes:** After each lecture, briefly summarize the material and extract its essence.
- **Work on Practice Problems:** Attempt to solve the practice problems and review my solutions at the end of the book.
- **Attend Office Hours:** Interact with the instructor, TAs, and other students.
- **Start Early:** Begin each problem set on the day it is released (or as soon as possible after that).
- **Finish Early:** Aim to complete problem sets and the midterm exam at least one day before the deadline.
- **Stuck? Ask for Help:** If you get stuck on a problem set problem, ask for hints on Piazza.

### Keep in Mind

My goal is to help you understand and learn the material. Understanding is a creative process that takes time and effort.

If you do not understand something, please ask me. If you are struggling to balance the workload, talk to me. If you have any concerns, let me know. Keep in mind that I am here to help.

### Honor Code

*“No member of the Caltech community shall take unfair advantage of any other member of the Caltech community.”*

---

<sup>3</sup> Based on Stanislav Dehaene, *How We Learn: The New Science of Education and the Brain*. 2020: Penguin Books Limited.