

Fall 2026

WEEK 11 STUDY GUIDE

The Big Picture

After two main applications of the mgf, we move to estimation from the frequentist and Bayesian perspectives

- We use the moment generating function to indicate why the CLT is true.
- The mgf and *Chernoff's bound* improves on the tail bounds of Markov and Chebyshev.
- In the frequentist world, a parameter is a fixed but possibly unknown number. The method of *maximum likelihood* identifies the parameter that makes the data most likely.
- In the Bayesian world, the data scientist's degree of uncertainty about unknown quantities is described by probability distributions. Unknown parameters are therefore random variables, and inference consists of updating our distribution of the parameter based on the observed data. For these calculations we need some basic methods for conditioning on continuous variables.
- Our main example is inference for the unknown p of a coin. Independence is affected by the randomization of the parameter.
- If the prior is uniform, the MAP estimate (the mode of the posterior) is the same as the maximum likelihood estimate with which the week started.

Week At a Glance

Mon 11/2	Tue 11/3	Wed 11/4	Thu 11/5	Fri 11/6
	Lecture	Sections	Lecture	Mega Section
HW 10 DUE AT NOON	HW 11 (Due 5PM Monday 11/9)			HW 11 Party 1 PM to 4 PM
No lab due	Lab 7A (Due 5PM Monday 11/9)		Lab 7A Party 4 PM to 6 PM	
MIDTERM 2 (8-10PM)	Work through Section 20.1	Work through Chapter 20; work through any remaining parts of Chapter 19	Work through Sections 21.1, 21.2	Skim Sections 22.1-22.2

Reading, Practice, and Class Meetings

Book	Topic	Lectures: Professor	Sections: TAs	Optional Additional Practice
Ch 20	**Approaches to inference** - 20.1 is about the method of maximum likelihood - 20.2 introduces conditioning on a continuous variable, and shows once again that randomizing a parameter affects dependence and independence (you saw this earlier with Poissonization) - 20.3 formalizes the concepts of prior and posterior distributions of parameters, and compares the MAP estimate and the MLE	Tuesday 11/3 - Proving CLT - Maximum likelihood - Random parameters: conditioning and independence	Wednesday 11/4 - Selections from HW 11	Ch 20 - Ex 2, 5
Ch 21	**Inference for the *p* of a random coin** - 21.1 picks up from 20.3, with a general beta prior instead of uniform - 21.2 is about the unconditional distribution of the number of heads, which is called beta-binomial - 21.3 is omitted this term	Thursday 11/5 - Inference for the random p of a coin - Conjugate priors; prediction - Relations between the beta and the binomial, including the beta-binomial distribution	Friday 11/6 - Ch 20 Ex 6 - Ch 21 Ex 2, 4	Ch 21 - All exercises not completed in section or homework