Probability Theory II
Lecture 1

Mon Jan/21/24.

Today

- Logistics
- What is this course about?
- Some examples?

Logistics

↳ Website

All the materials available here.

mateodd25.github.io/prob2

↳ Instructor

Use Piazza before emailing me!

Mateo Díaz (mateodd@jhu.edu)
OH: Monday 3-5pm Wyman S429

↳ TA

Ao Sun
OH: Friday 9-11 am Wyman S425

↳ Textbook

See link on the website.

Probabity: Theory & Examples. Rick Durrett.

↳ Grading System

Grade will have 4 components:

- Homework (40%) ← Every ~2 weeks.

- Take-home Exam (25%) ← ~ Midterm.
- Final project (25%) ← Document + Presentation
- Participation (10%) ← Ask questions!

↳ Submitting work

Everything will be submitted via Gradescope.

What is this course about?

So far what in Probability Theory 1 we covered measure theoretic foundations. But most of it was "static." This semester we will cover dynamic stochastic processes.

- Characteristic functions.
- Poisson process.
- Conditional probability and expectation.
- Martingales. ← The CLT that we missed.
- Kolmogorov's extension Theorem.
- Markov's chains.
- Random Walks in $\mathbb{R}^n$.
- Brownian Motion.

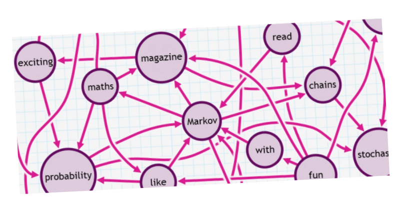

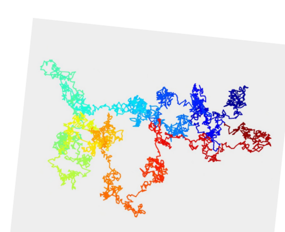

If we have enough time we will cover

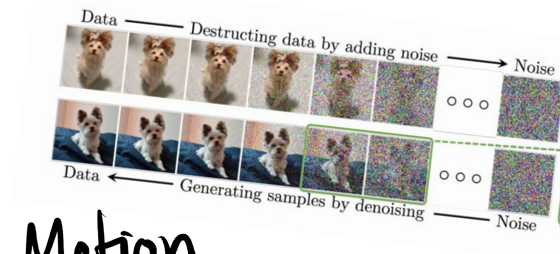

- Ergodic Theory
- Stochastic Calculus.
- Multidimensional Brownian Motion.

Some examples

- Monkeys writing Shakespeare.

Imagine you have a Monkey randomly typing letters in a computer.

Assume that they type any letter uniformly at random.

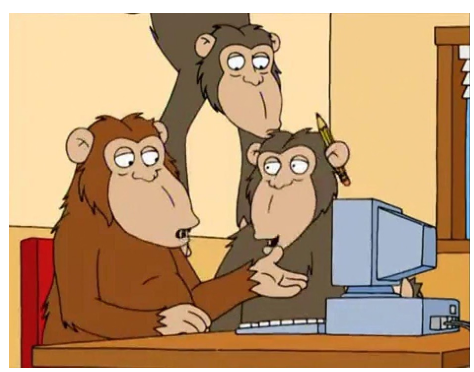

Any finite sequence of letters has positive probability of appearing. So the monkey will eventually type:

- A Shakespeare Novel.

- Your DNA code,
- Your phone number,
- The solutions to our homework problems.

⋮

Question: On average how long would it take the Monkey to write "abracadabra"?

We will tackle this problem using martingales. Can you think of a way of doing it from first principles?

- Stochastic algorithms.

Often in large scale optimization we are interested in solving

$$\min_{x} \mathbb{E}_z f(x; z)$$

Random variable

decision variable

We often do not know the distribution of z and only have access.

Often, people solve these problem via SGD:

$$x_{k+1} \leftarrow x_k - \overset{\text{stepsize}}{\alpha_k} \underbrace{\nabla f(x_k, z_k)}_{\text{stochastic approximation of the gradient.}}$$

Since z_k is random, then x_{k+1} is random.

Illustrative example: Imagine you are Uber/Lyft, you want to solve:

$$\max \mathbb{E}_z \underset{\text{Revenue}}{f}(\overset{\text{Pricing}}{x}, \underset{\text{Demand}}{z})$$

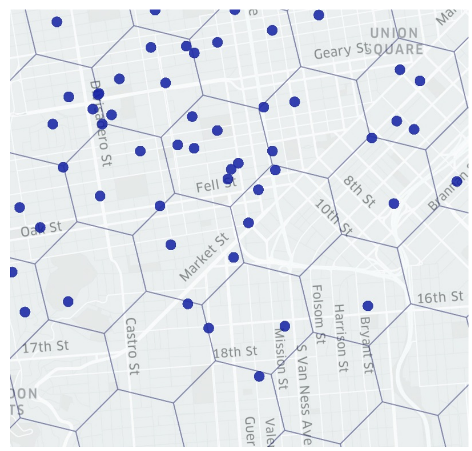

Conditional expectation

Question. When does x_k converge to a solution? If it converges, can we have confidence region? Martingales!

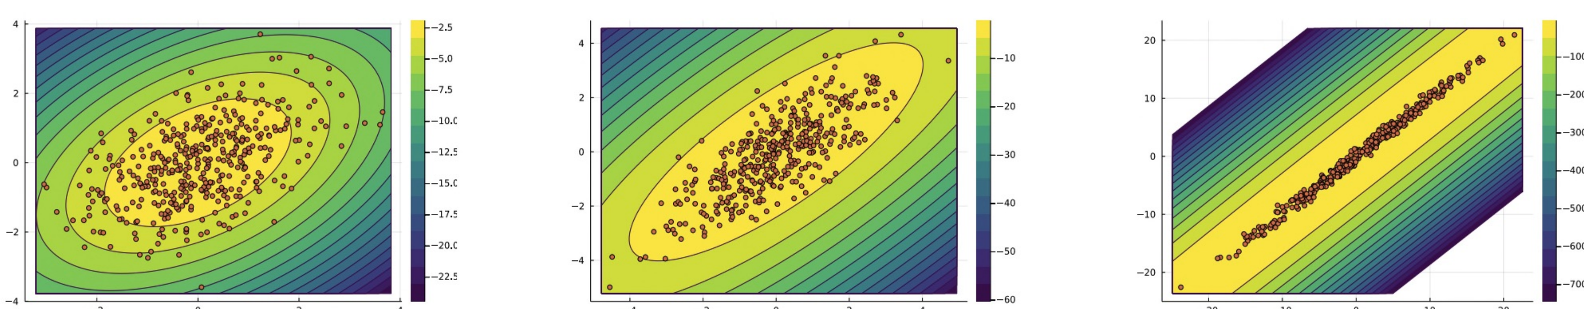

▷ Shuffling cards

Assume you work at a cheap casino and have to shuffle any new desk of cards at the start of a game.

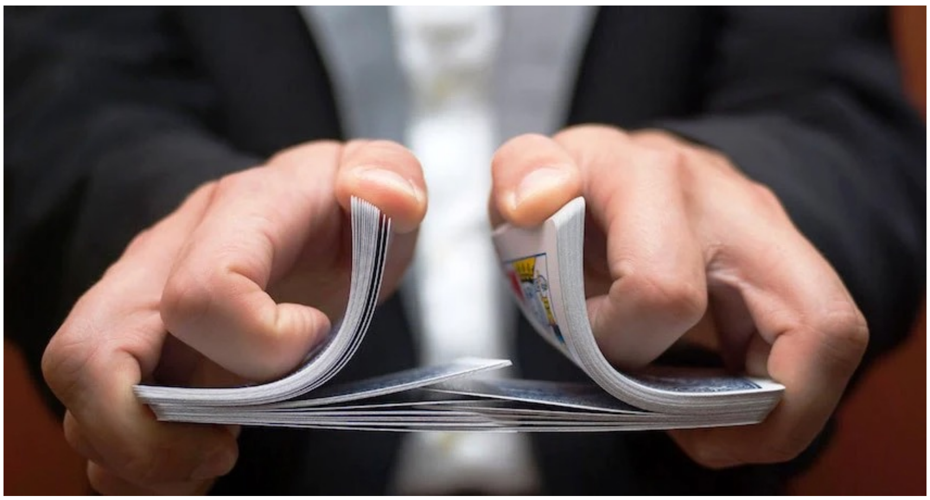

Question. How many times do you need to shuffle the deck before it looks random?

We can tackle this question with Markov Chains (you need to do it 8 times).

▷ Drunk humans vs Drunk birds

A drunk grad student walks randomly

on a grid in 2D. Simultaneously a drunk bird flies randomly in a 3D grid.

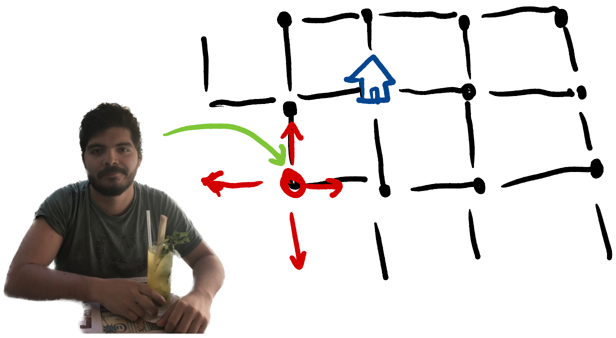

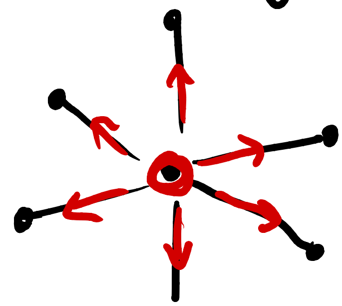

Question: Is the grad student guaranteed to eventually make it home? How about the bird?

We will tackle this question once we cover random walks in $\mathbb{R}^d$.

▷ Particles suspended in a fluid

Particles are constantly moving and colliding with each other.

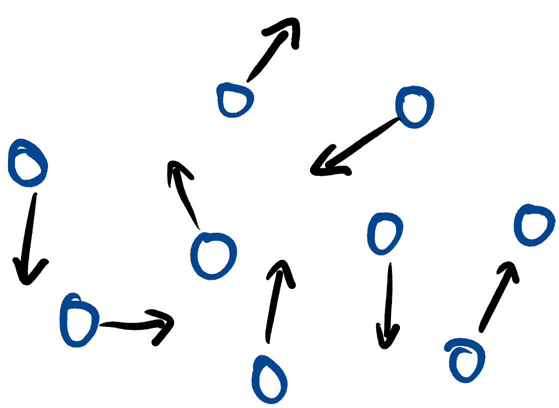

In 1827, Robert Brown was studying grain of pollen in water and he noticed a continuous jittery motion.

Question: How can we model the random movement of these particles?

Brownian Motion!